Online Algorithms for Information Aggregation from Distributed and Correlated Sources

Chi-Kin Chau, *Member*, *IEEE*, Majid Khonji, and Muhammad Aftab

arXiv:1601.03147v2 [cs.NI] 13 Apr 2016

***Abstract*—There is a fundamental trade-off between the communication cost and latency in information aggregation. Aggregating multiple communication messages over time can alleviate overhead and improve energy efficiency on one hand, but inevitably incurs information delay on the other hand. In the presence of uncertain future inputs, this trade-off should be balanced in an online manner, which is studied by the classical *dynamic TCP ACK problem* for a single information source. In this paper, we extend dynamic TCP ACK problem to a general setting of collecting aggregate information from distributed and correlated information sources. In this model, distributed sources observe correlated events, whereas only a small number of reports are required from the sources. The sources make online decisions about their reporting operations in a distributed manner without prior knowledge of the local observations at others. Our problem captures a wide range of applications, such as in-situ sensing, anycast acknowledgement and distributed caching. We present simple threshold-based competitive *distributed online algorithms* under different settings of intercommunication. Our algorithms match the theoretical lower bounds in order of magnitude. We observe that our algorithms can produce satisfactory performance in simulations and practical testbed.**

***Index Terms*—Online algorithms, competitive analysis, distributed optimization problem, latency, communication cost**

I. Introduction

Information often arises from multiple sources, and efficient mechanisms are required to aggregate the information from these sources. One can aggregate multiple communication messages to alleviate overhead and improve energy efficiency. However, there is a fundamental trade-off between the communication cost and latency in information aggregation. Aggregating multiple communication messages over time can reduce the communication cost on one hand, but inevitably incurs information delay on the other hand. Typically, there is uncertainty in future inputs, and the balance between communication cost and latency needs to be tackled in an online manner. A well-known result is the classical dynamic TCP ACK problem [1], which studies the balance between communication cost (i.e., the number of ACKs) and latency (i.e., the lag between packet reception and ACK transmission) in an online manner for a single information source.

Notably, there are similar trade-off balancing problems in a wide range of information exchange applications, especially when there are distributed systems sharing correlated local information. A critical problem is how to coordinate the operations of these systems for jointly balancing the communication cost and latency of information exchanges. There are several concrete examples as follows:

- *In-situ Sensing*: A set of sensors are assigned to monitor certain events of interest in their vicinity based on local measurements. The measurements are correlated and are required to report to a base station. The communication messages from sensors should be optimized to reduce battery consumption and improve operation lifetime, while incurring tolerable latency of reporting.
- *Anycast ACKs*: A sender is transmitting a number of packets to a group of anycast receivers in an unreliable channel. An ACK from one of the successful receivers is required to confirm the reception. Furthermore, the receiver may reply an ACK to a batch of packets, instead of to every packet. To improve throughput, the number of ACKs needs to be minimized among the receivers, taking into account the incurred latency.
- *Distributed Caching*: For distributed storage and web caching services, the cached objects in data servers are often represented by a hash file, which is distributed among other servers for coordinating information retrieval. The total size of hash files generated by all servers should be optimized, by taking into account the commonly cached objects and the latency of distributing these hash files.

The aforementioned applications often feature a trade-off between a certain communication cost (e.g., energy consumption, bandwidth and storage space) and the incurred latency in an online fashion, similar to that of dynamic TCP ACK problem. However, as a departure from the classical TCP ACK problem, these problems involve multiple information sources, and the difficulty is that the local information of each source is typically unknown to others without communication, which hinders the joint optimization. We consider a general problem of a general setting of collecting aggregate information from distributed information sources as follows:

- *Online Decisions:* A group of systems are tasked with reporting certain correlated events (e.g., sensing events, receptions of packets, and cached objects). We assume no prior knowledge on the arrivals of events (either stochastic or non-stochastic models). The systems are required to make decisions in an online manner.
- *Distributed Decisions:* The local decisions of the systems about their report messages should globally optimize the communication cost and the incurred latency. In this paper, we consider various settings: (i) *no intercommunication*, where each system is not aware of the communications of others, (ii) *full intercommunication*,

C.-K. Chau, M. Khonji, and M. Aftab are with the Department of EECS at Masdar Institute of Science and Technology, UAE (e-mail: {ckchau, mkhonji, muhaftab}@masdar.ac.ae).

where every system can overhear the communications of others, and (iii) *partial intercommunication*, such that the extent of intercommunication is characterized by a communication graph.

Not only generalizing the dynamic TCP ACK problem to a setting with distributed information sources, our problem is also generalized to consider general communication cost that captures properties, such as energy consumption, bandwidth or storage space, and general latency penalty that models considerations, such as the ineffectiveness of delayed information, dissatisfaction of users, and incurred economic loss. Our problem can be applied to a broad range of applications.

In the literature [2], *competitive ratio* is a common metric to compare the performance of an online algorithm against an optimal offline algorithm. In our context, we compare the performance of a *distributed* online algorithm against a *centralized* optimal offline algorithm. We provide simple threshold-based competitive distributed online algorithms in various settings for our problem that match with the theoretical lower bounds of competitive ratios in order of magnitude. Moreover, we observe that our algorithms can produce satisfactory performance in both simulations and practical testbed.

Outline: The problem is formulated in Sec. III. Simple online algorithms for various intercommunication settings are presented in Secs. V-VI, which match with the lower bounds of competitive ratios in Sec. VII. We performed simulations to corroborate the performance in Sec. VIII. We also implemented our algorithms in a practical sensor network testbed, with observations discussed in Sec. IX. Some proofs are deferred to the technical report [3], unless stated otherwise.

II. Related Work

A. Online Algorithms

Online algorithms are critical for a wide range of problems [2], [4], [5] involving uncertain input and timely decisions, by making decisions and generating outputs instantaneously over time, based on only the part of the input that have seen so far, without knowing the rest in the future. Online algorithms have several benefits. First, they do not require a-prior or stochastic knowledge of the inputs, which is robust in any uncertain (even adversarial) environments, such as in dynamic ad hoc networks. Second, online algorithms are often simple decision-making mechanisms, without relying on future information prediction. Third, online algorithms give a fundamental characterization of the problems, which are useful to benchmark more sophisticated decision-making mechanisms.

In this paper, we study a *distributed* online algorithmic problem, with multiple online decision-makers. This paper also extends the online decision problem in [6], which considers a single sensor for temperature control in an online fashion.

B. Dynamic TCP Acknowledgement

Dynamic TCP ACK problem [1], [7], [8] is equivalent to a subclass of our problem (i.e., the one-system case). There are a stream of packets arriving at a destination. The packets must be acknowledged in order to notify the sender that the transmission was successful. However, it is possible to simultaneously acknowledge multiple packets using a single ACK packet. The delayed acknowledgement reduces the number of the ACKs, but might incur excessive latency to the TCP connection and interfere with the TCP's congestion control mechanism. The problem is to find a solution with the minimum cost of the total number of ACKs sent and the latency cost incurred by delayed ACKs. Dooly et al. [1] formulated this problem, and devised a 2-competitive deterministic algorithm. Dynamic TCP ACK problem may be viewed as a recurrent version of ski-rental problem [9]. Despite of many variants of ski-rental problem in the literature [8], [10], there lack results about distributed online decision-makers.

C. Information Aggregation and Multicast Acknowledgement

Information aggregation is a useful technique for in-situ sensing, which has been studied extensively in the literature [11]. The prior literature considered various system solutions, such as MAC and routing protocols that incorporate aggregation, aggregation functions that merge information from multiple sources, and representations of data to facilitate aggregation. In contrast, this paper considers the online decision-making mechanisms for information aggregation, which may be employed in tandem with the prior system solutions.

On the other hand, a closely related problem to our work is Multicast Acknowledgement Problem (MAP) [12], [13]. In MAP, a root sends data packets along a multicast tree, which has unreliable transmission links. Each node in the multicast tree are required to reply ACKs to the root for its received data packets. Rather than transmitting an ACK immediately, each node can delay its ACKs, by aggregating the ACKs for multiple data packets and the ACKs from the descendant nodes in the multicast tree. The key differences of MAP with our problem are that we do not require an ACK from every node (but only K ACKs from all the receivers), and we consider non-hierarchical network such that every node can communicate with base station (unlike the tree hierarchy in MAP). The differences lead to different characterizations of competitive ratios of online algorithms. We also note that there is a recent extension of MAP to consider bi-criteria optimization of communication cost and latency separately, rather than by a weighted objective function [14], [15].

III. Problem Formulation

A. Main Components

This paper considers a general model of a group of distributed systems, which make online decisions for reporting certain correlatively observed events in a distributed manner.

The model is consisted of the following components:

- **Systems**: A group of N systems, indexed by set $\mathcal{N}$, are tasked with reporting certain events to a base station.
- **Events**: A sequence of events of interest, indexed by a finite set $\mathcal{M}$, appear over a continuous time horizon. For each event $j \in \mathcal{M}$, let its appearance time be $\mathtt{t}^j$. To model the observation by a system, let w_i^j represent an

abstract real-valued metric of observation for event j at system i. If event j is not observed by i, then $w_i^j = 0$. Denote the tuple of measurement of j at i by $\sigma_i^j \triangleq (\mathtt{t}^j, w_i^j)$, and the measurements at all systems by $\sigma \triangleq (\sigma_i^j)_{i\in\mathcal{N},j\in\mathcal{M}}$. We assume that there is only one event occurred at each t, because of continuous time.

- **Reports**: For each event, at least $K(\leq N)$ report messages from K systems in $\mathcal{N}$ that observe the event (i.e., $w_i^j > 0$) are required by the base station. General content of a report message is allowed. For instance, a report message for notification may only include a Boolean flag if an event is observed. A report message may also contain more detailed information, such as the appearance times and measurement values of the observed events. This paper only considers the abstract communication cost associated with the report messages.

Example 1: There are a few examples of the model:

- **Anycast ACKs**: An event is a reception of a packet j at an anycast receiver i. This is an extension of the classical (unicast) TCP ACK problem to a distributed (anycast) setting. Let $K = 1$, and $w_i^j = 1$ if packet j is received by receiver i. A report message is an ACK packet.
- **In-situ Sensing**: An event is an observation of a substance source j at a sensor i, and w_i^j is the observed concentration level of source j at sensor i. For example, in [6], w_i^j is the temperature fluctuation, when a heat source appears. A report message is a measurement report from a sensor. Sometimes, multiple report messages are required. For instance, tracing the location of the source may require the measurements from $K \geq 3$ sensors.
- **Distributed Caching**: An event is an object j being cached at a distributed server i, and w_i^j represents the size of object j, as cached by server i. A report message is a hash file of the cached objects from a server. The hash files will be distributed among other servers for notifying the presence of cached objects at other servers. For reliability, it is desirable to cache each object with at least K servers, if available.

B. Problem Definition

The systems in $\mathcal{N}$ make decisions whether to report the observed events at time t. Let the decisions of system i be a sequence $\mathcal{R}_i = \big(r_{i,k}, (x_k(i,j))_{j\in\mathcal{M}}\big)_{k=1}^{\eta_i}$, where $r_{i,k}$ is the reporting time of the k-th report message, and $x_k(i,j) \in \{0,1\}$ is an indicator variable whether measurement σ_i^j is included in the k-th report message. Let η_i be the total number of report messages from i. Let $\gamma^j(K)$ be the reporting time of the K-th report[1] from one of the observed systems for event j. Namely,

$$\gamma^j(K) \triangleq K^{\text{th}}\text{-smallest}\{r_{i,k} \mid x_k(i,j) = 1 \wedge w_i^j > 0\}$$

where K^{th}-smallest selects the K-th smallest item from a set of numerical items, with a deterministic tie-breaking rule.

Denote the set of measurements included in the k-th report message of system i by $\sigma_{i,k} \triangleq \{\sigma_i^j : x_k(i,j) = 1\}$.

[1] If event j is never reported for K times, then $\gamma^j(K) = \infty$. Without loss of generality, it is assumed that each event must trigger non-zero measurement values for at least K systems.

The goal of K-report **D**istributed **I**nformation **A**ggregation problem (K-DIA) is to minimize a global objective function of the communication cost of all report messages and the associated latency penalty for the K-th report messages, as defined as follows:

$$(K\text{-DIA}) \qquad \min_{\mathcal{R}_i: i\in\mathcal{N}} \sum_{i\in\mathcal{N}} \Big(\rho \cdot \sum_{k=1}^{\eta_i} \mathcal{C}(\sigma_{i,k}) + (1-\rho)\cdot \sum_{j\in\mathcal{M}} \mathcal{L}(\sigma_i^j, \gamma^j(K))\Big) \tag{1}$$

$$\text{s.t.} \qquad \mathtt{t}^j \leq r_{i,k} \text{ if } x_k(i,j) = 1, \text{ for all } j \in \mathcal{M} \tag{2}$$

The notations in the objective are explained as follows:

1) $\mathcal{C}(\sigma_{i,k})$ represents the abstract *communication cost* of a report message containing the set of measurements $\sigma_{i,k}$. The communication cost can reflect various properties, such as energy consumption, bandwidth or storage space of the report messages. We consider a general communication cost function. For example, if the communication cost depends on only the number of messages (e.g., for notification), then $\mathcal{C}(\sigma_{i,k}) = 1$. If the communication cost depends on the binary representation of the total measurement (e.g., for data archiving), then
$$\mathcal{C}(\sigma_{i,k}) = C_1 \cdot \log\Big(\big[\sum_{j\in\mathcal{M}:\sigma_i^j\in\sigma_{i,k}} w_i^j\big]_{\underline{w}}^{\bar{w}}\Big) + C_0 \tag{3}$$
where $[\cdot]_{\underline{w}}^{\bar{w}} \triangleq \min(\max(\cdot, \underline{w}), \bar{w})$ is an operator that rounds the measurement value within the range $[\underline{w}, \bar{w}]$. Generally, $\mathcal{C}(\sigma_{i,k})$ is assumed to be non-decreasing in each parameter w_i^j, and *sub-additive*, such that
$$\mathcal{C}(\sigma_{i,k} \cup \sigma'_{i,k}) \leq \mathcal{C}(\sigma_{i,k}) + \mathcal{C}(\sigma'_{i,k}) \tag{4}$$
Sub-additivity reflects the benefit of overhead reduction by aggregating multiple report messages into a single one. Suppose the communication cost per report is bounded by $\mathcal{C}_{\min} \leq \mathcal{C}(\sigma_{i,k}) \leq \mathcal{C}_{\max}$. Let $\alpha \triangleq \frac{\mathcal{C}_{\max}}{\mathcal{C}_{\min}}$, where $\alpha \geq 1$.
2) $\mathcal{L}(\sigma_i^j, \gamma^j(K))$ represents the *latency penalty* that captures the adverse effect when the K-th report message of event j is delayed until $\gamma^j(K)$. The latency penalty can model factors, such as the ineffectiveness of delayed information, dissatisfaction of users, and incurred economic loss. We consider a general penalty function. For example, the penalty function can be a linear function:
$$\mathcal{L}(\sigma_i^j, \gamma^j(K)) = w_i^j \cdot (\gamma^j(K) - \mathtt{t}^j) \tag{5}$$
Generally, $\mathcal{L}(\sigma_i^j, \gamma^j(K))$ is assumed to be non-decreasing in each w_i^j and $\gamma^j(K)$, and $\mathcal{L}(\sigma_i^j, \mathtt{t}^j) = 0$ (because there is no delay) and $\mathcal{L}(\sigma_i^j, \gamma^j(K)) = 0$ if $w_i^j = 0$ (because the event is not observed).
3) ρ is a parameter for controlling the trade-off between the communication cost and latency. A larger value of ρ favors the reduction in communication cost, whereas a smaller value of ρ favors the reduction in latency.

Note that there is a trade-off in information aggregation between the communication cost and latency penalty. K-DIA

is a generalization of the classical dynamic TCP ACK problem [1] by considering information aggregation from multiple sources.

For the convenience of illustration, we also define a sub-problem called 1-sDIA, which particularly captures the anycast ACK problem:

$$\text{(1-sDIA)} \qquad \min_{\mathcal{R}_i: i\in\mathcal{N}} \sum_{i\in\mathcal{N}} \Big(\rho \cdot \sum_{k=1}^{\eta_i} \mathcal{C}(\sigma_{i,k}) + (1-\rho)\cdot \sum_{j\in\mathcal{M}} \mathcal{L}(\sigma_i^j, \gamma^j(1))\Big) \tag{6}$$

where $\mathcal{C}(\sigma_{i,k}) = 1$ and $\mathcal{L}(\sigma_i^j, \gamma^j(1)) = w_i^j \cdot (\gamma^j(1) - \mathtt{t}^j)$, and

$$\gamma^j(1) = \min_{i\in\mathcal{N}, k\in\{1,\ldots,\eta_i\}} \{r_{i,k} \mid r_{i,k} \geq \mathtt{t}^j \wedge w_i^j > 0\}$$

Note that the classical TCP ACK problem is the one-system case (i.e., $|\mathcal{N}| = 1$) of 1-sDIA.

C. Distributed Online Algorithms

Given all the inputs $\sigma = (\mathtt{t}^j, (w_i^j)_{i\in\mathcal{N}})_{j\in\mathcal{M}}$ in advance, K-DIA can be solved in an offline manner. However, the measurements σ are revealed gradually over time in practice. Furthermore, in multi-system setting, each system initially only knows its local measurements $(\sigma_i^j)_{j\in\mathcal{M}}$, without prior knowledge of the local observations at others. Thus, there is incomplete information at each system about the inputs. The incomplete information presents a unique challenge in the multi-system setting, as compared to the one-system setting (i.e., dynamic TCP ACK problem).

We will provide *distributed online* algorithms for solving K-DIA. More specifically, we focus on *deterministic* distributed algorithms with the following characteristics:

1) The decisions $\mathcal{R}_i$ at each system follow a set of pre-defined deterministic rules.
2) The measurements $(\sigma_i^j)_{j\in\mathcal{M}}$ are observed locally by each system i only, and cannot be revealed to other systems without involving communication.
3) If there involves any communication among the systems, then this also incurs a cost per transmission, which may be as *costly* as communicating a report message to the base station. Thus, it may be more efficient by communicating a report message to the base station, and letting other systems overhear it, if overhearing is possible.

An algorithm is called *online*, if the decision of at the time t only depends on the information available before or at t. Namely, the decision of producing a report message at time t is only based on the measurements $(\sigma_i^j)_{j:\mathtt{t}^j\leq t, i\in\mathcal{N}}$.

A simple class of online algorithms are threshold-based algorithms. First, define the *accumulative latency penalty* at system i for an interval $(t_1, t_2]$ by

$$\mathtt{lat}_i(t_1, t_2) \triangleq \sum_{j\in\mathcal{M}:\mathtt{t}^j\geq t_1} \mathcal{L}(\sigma_i^j, t_2) \tag{7}$$

and the *accumulative communication cost* by

$$\mathtt{com}_i(t_1, t_2) \triangleq \mathcal{C}\big(\{\sigma_i^j : t_1 < \mathtt{t}^j \leq t_2\}\big) \tag{8}$$

In a *threshold-based* algorithm, each report message includes all the unreported events that are observed locally. A report message is produced, if the ratio between the accumulative latency penalty and the accumulative communication cost since the last report message exceeds a pre-defined threshold θ. For example, in 1-DIA, the reporting time of the k-th report message at system i is obtained by

$$r_{i,k} = \min\Big\{t : \frac{\mathtt{lat}_i(r_{i,k-1}, t)}{\mathtt{com}_i(r_{i,k-1}, t)} \geq \theta\Big\} \tag{9}$$

Threshold-based algorithm is online, since the accumulative latency penalty and accumulative communication cost depend on no future information. This paper studies the effectiveness of threshold-based algorithms.

D. Intercommunication

In addition to the report messages to the base station, extra intercommunication may be enabled among systems. We consider implicit intercommunication, such that one system may overhear the messages that are not intended for it (e.g., other report messages to the base station). The extent of intercommunication is characterized by the following cases:

1) **No Intercommunication:** There is no intercommunication among the systems. Namely, no system needs to communicate among themselves, nor can they overhear from each other. But the systems can communicate to the base station for reporting only. For example, in anycast over a wireline network, overhearing is difficult.
2) **Full Intercommunication:** In contrast to no intercommunication, there is full intercommunication among the systems. For instance, a system can overhear report messages produced by any other systems and hence, possibly withholds its own report messages. This can be achieved, for instance, in a broadcast channel or by replicating the messages by the base station to other systems. For example, in wireless sensor networks, overhearing is possible. An ideal setting of full intercommunication is *collision-free*, if no pair of systems transmits messages simultaneously[2]. This can effectively eliminate redundant report messages produced at the same time.
3) **Partial Intercommunication:** In practice, the extent of intercommunication is not prefect, which lies between the two preceding cases. We model the availability of intercommunication among systems by a *communication graph*, such that only the neighboring systems in the communication graph can engage in full intercommunication (including overhearing) with each other.

In the subsequent sections, K-DIA will be studied under these cases of intercommunication.

Example 2: This example considers 1-sDIA. The basic idea of the problem and the operation of threshold-based algorithm

[2] *Collision-free* intercommunication can be implemented by several ways. For example: (1) Priority based: Given a pre-defined priority order, a system in lower priority is required to give way to those systems in higher priority by spending more time to sense the channel before transmission. (2) Random access: Before transmission, each system should wait for an exponentially distributed random interval before transmission. The probability of collision is negligible in continuous time.

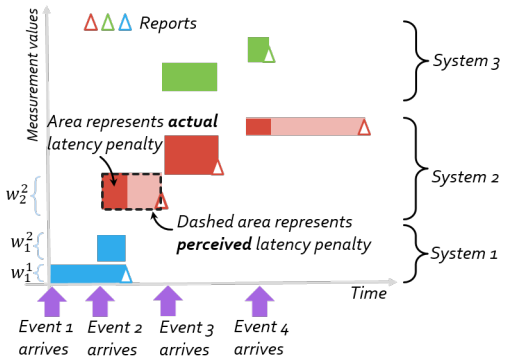

Fig. 1: An illustration of threshold-based algorithm for 1-sDIA. Four events arrive over time, whose measurement values to each system are plotted in y-axis. The rectangles represent the latency penalty, whereas the triangles represent the report messages.

are illustrated in Fig. 1. In this example, system 2 cannot overhear others, but other systems can. Note that the communication cost in 1-sDIA is always one per message. The threshold-based algorithm is applied as follows:

1) All systems defer the report messages until the perceived accumulative latency penalty since the last report message exceeds a threshold. Each report message includes all unreported events that are observed locally.
2) If a system overhears a report message produced by another system, it needs not transmit a report message for those reported events. It resets the accumulative latency penalty, if all observed events have been reported.

In Fig. 1, the actual accumulative latency penalty of each event is represented by the rectangles in deep colors. Since system 2 is not aware of the report messages produced by others, its perceived accumulative latency penalty may be more than the actual one, which is represented by the total area of deep and light colors (e.g., the rectangle in dashed line). In this example, five report messages are generated.

E. Competitive Analysis and Competitive Ratio

In the literature of online algorithms, competitive analysis is a common framework to characterize the performance [2].

Denote the cost of online algorithm $\mathcal{A}$ by $\texttt{Cost}(\mathcal{A}[\sigma])$ and the cost of an optimal offline solution on input σ by $\texttt{Opt}(\sigma)$. The *competitive ratio* for $\mathcal{A}$ is defined as the worst-case ratio between the cost of the online algorithm and that of an optimal offline solution, namely,

$$\texttt{CR}(\mathcal{A}) \triangleq \max_{\sigma} \frac{\texttt{Cost}(\mathcal{A}[\sigma])}{\texttt{Opt}(\sigma)} \tag{10}$$

$\mathcal{A}$ is called c-competitive, if $\texttt{CR}(\mathcal{A}) = c$.

In the context of distributed online algorithms, we compare with the (centralized) optimal offline solution, which knows all the inputs in advance without the need for intercommunication.

In this paper, we show that threshold-based algorithms with proper thresholds can produce competitive performance for solving K-DIA that are close to the theoretical lower bounds of any deterministic online algorithms. The main results are summarized in Table I.

Competitive Ratio	No Intercommunication	Full Intercommunication
Threshold-based	$\frac{\alpha N}{K}+1$	$\frac{\sqrt{(K-\alpha)^2+4\alpha K N}+\alpha+K}{2K}$
Lower bound ($K=1$)	$\alpha(N-1)+\frac{3}{2}$	$\frac{\sqrt{N}}{4}$
Competitive Ratio	**Partial Intercommunication**	
Threshold-based	$\frac{\sqrt{(\alpha \mathtt{x}-K)^2+4\alpha K N}+\alpha \mathtt{x}+K}{2K}$ where $1 \le \mathtt{x} \le N$ is a network parameter	

TABLE I: Competitive ratios of threshold-based algorithms, and the theoretical lower bounds of any deterministic online algorithms.

IV. No Intercommunication

This section considers the simplest setting with no intercommunication among the systems. Threshold-based algorithm $\mathcal{A}_{\text{thb}}$ (Algorithm 1) produces a report message independently at each system i when the ratio between the perceived accumulative latency penalty and the accumulative communication cost since the last report message exceeds threshold θ. Note that for each event, $\mathcal{A}_{\text{thb}}$ produces report messages at all the systems that observe non-zero measurement values. Hence, every event triggers $N(\ge K)$ report messages, and this produces a feasible solution for K-DIA.

Algorithm 1 $\mathcal{A}_{\text{thb}}[\theta, t_{\text{now}}, (\sigma_i^j)_{j:\mathtt{t}^j \le t_{\text{now}}}]$

Input: Threshold θ; current time t_{now}; known measurements $(\sigma_i^j)_{j:\mathtt{t}^j \le t_{\text{now}}}$
Global Init.: $x_k(i,j) \leftarrow 0$ for all $j \in \mathcal{M}, k \leftarrow 0, r_{i,k} \leftarrow 0$

1: **if** $\frac{\texttt{lat}_i(r_{i,k}, t_{\text{now}})}{\texttt{com}_i(r_{i,k}, t_{\text{now}})} \ge \theta$ **then** ▷ *Condition on threshold* θ
2: $k \leftarrow k+1$ ▷ *Next report message*
3: $x_k(i,j) \leftarrow 1$ for all observed event j
4: $r_{i,k} \leftarrow t_{\text{now}}$
5: **end if**

Denote by $\mathcal{A}_{\text{thb}}[\theta, \sigma]$ the threshold-based algorithm with threshold θ on input σ. By a slight abuse of notation, we also denote by $\texttt{Opt}(\sigma)$ the optimal offline solution on input σ.

First, we consider 1-DIA of at least one report message per event. Given input σ, let $x_k^*(j) \in \{0,1\}$ be the indicator variable whether event j is included in the k-th report message by $\texttt{Opt}(\sigma)$. The set of events $\mathcal{M}$ can be partitioned into a set of segments $\mathcal{S}^*$, grouped by the corresponding report messages, such that $\mathtt{S}_k \in \mathcal{S}^*$ is defined by

$$\mathtt{S}_k \triangleq \{j \in \mathcal{M} \mid x_k^*(j) = 1\} \tag{11}$$

For each segment $\mathtt{S} \in \mathcal{S}^*$, denote the sub-sequence of measurements in $\mathtt{S}$ by $\sigma_{\mathtt{S}} \triangleq (\sigma_i^j)_{i \in \mathcal{N}, j \in \mathtt{S}}$. The competitive ratio can be obtained considering each $\sigma_{\mathtt{S}}$.

Theorem 1: For 1-DIA, let threshold $\theta = \frac{\rho}{\alpha N(1-\rho)}$ for every system $i \in \mathcal{N}$, the competitive ratio of $\mathcal{A}_{\text{thb}}$ is $\texttt{CR}(\mathcal{A}_{\text{thb}}[\frac{\rho}{\alpha N(1-\rho)}]) = \alpha N + 1$.

Proof: The optimal offline solution $\texttt{Opt}(\sigma)$ determines a set of segments $\mathcal{S}^*$. Note that $\texttt{Opt}(\sigma) = \sum_{\mathtt{S} \in \mathcal{S}^*} \texttt{Opt}(\sigma_{\mathtt{S}})$. By Lemma 1 (see [3]), we obtain

$$\texttt{Cost}(\mathcal{A}_{\text{thb}}[\theta, \sigma]) \le \sum_{\mathtt{S} \in \mathcal{S}^*} \texttt{Cost}(\mathcal{A}_{\text{thb}}[\theta, \sigma_{\mathtt{S}}]) \tag{12}$$

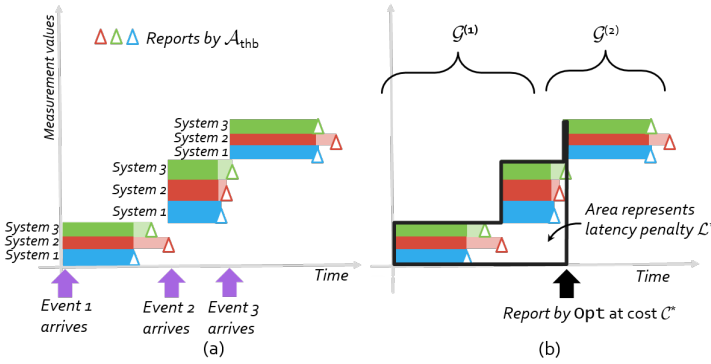

Fig. 2: Consider 1-sDIA. (a) $\mathcal{A}_{\text{thb}}$ produces a report message independently at each system. (b) Opt produces a report message in the end of each segment $\mathtt{S} \in \mathcal{S}^*$ with communication cost $\mathcal{C}^*$ and latency penalty $\mathcal{L}^*$ (i.e., the polygon area in black border). The set of measurements is divided into two groups $\mathcal{G}^{(1)}$, $\mathcal{G}^{(2)}$.

Since $\frac{\sum_l a_l}{\sum_l b_l} \leq \max_l\{\frac{a_l}{b_l}\}$ for any positive a_l, b_l, we obtain

$$\frac{\mathtt{Cost}(\mathcal{A}_{\text{thb}}[\frac{\rho}{\alpha N(1-\rho)}, \sigma])}{\mathtt{Opt}(\sigma)} \leq \max_{\mathtt{S}\in\mathcal{S}^*} \frac{\mathtt{Cost}(\mathcal{A}_{\text{thb}}[\frac{\rho}{\alpha N(1-\rho)}, \sigma_{\mathtt{S}}])}{\mathtt{Opt}(\sigma_{\mathtt{S}})} \tag{13}$$

The rest of proof aims to find an upper bound for $\frac{\mathtt{Cost}(\mathcal{A}_{\text{thb}}[\frac{\rho}{\alpha N(1-\rho)}, \sigma_{\mathtt{S}}])}{\mathtt{Opt}(\sigma_{\mathtt{S}})}$ for any given S.

Recall that $\mathtt{Opt}(\sigma_{\mathtt{S}})$ produces a report message in the end of S. Let the optimal offline cost be $\mathtt{Opt}(\sigma_{\mathtt{S}}) = \rho\mathcal{C}^* + (1-\rho)\mathcal{L}^*$, where $\mathcal{C}^*$ and $\mathcal{L}^*$ are the communication cost and latency penalty of $\mathtt{Opt}(\sigma_{\mathtt{S}})$ respectively.

Next, we divide the set of measurements $\{\sigma_i^j\}_{j\in\mathtt{S}, i\in\mathcal{N}}$ into two groups $\mathcal{G}^{(1)}, \mathcal{G}^{(2)}$:

- $\mathcal{G}^{(1)}$ is the set of measurements in $\{\sigma_i^j\}_{j\in\mathtt{S}, i\in\mathcal{N}}$, such that $\mathcal{A}_{\text{thb}}[\frac{\rho}{\alpha N(1-\rho)}]$ produces report messages before or at the same time as $\mathtt{Opt}(\sigma_{\mathtt{S}})$ does.
- $\mathcal{G}^{(2)}$ is the set of measurements in $\{\sigma_i^j\}_{j\in\mathtt{S}, i\in\mathcal{N}}$, such that $\mathcal{A}_{\text{thb}}[\frac{\rho}{\alpha N(1-\rho)}]$ produces report messages after $\mathtt{Opt}(\sigma_{\mathtt{S}})$ does. Note that one system can at most produce one report message for $\mathcal{G}^{(2)}$.

See Fig. 2 for an illustration considering 1-sDIA.

Regarding to $\mathcal{A}_{\text{thb}}[\frac{\rho}{\alpha N(1-\rho)}]$, let the actual total latency penalty for $\mathcal{G}^{(1)}$ be $\mathcal{L}^{(1)}$, and the perceived total latency penalty be $\hat{\mathcal{L}}^{(1)}$. Note that $\mathcal{L}^{(1)} \leq \hat{\mathcal{L}}^{(1)}$ because the systems cannot overhear the report messages from other systems. Let the total communication cost for $\mathcal{G}^{(1)}$ be $\mathcal{C}^{(1)}$. Similarly, we define $\mathcal{L}^{(2)}$, $\hat{\mathcal{L}}^{(2)}$ and $\mathcal{C}^{(2)}$ for $\mathcal{G}^{(2)}$. Since $\mathcal{A}_{\text{thb}}[\frac{\rho}{\alpha N(1-\rho)}]$ is a threshold-based algorithm, it follows that

$$\hat{\mathcal{L}}^{(1)} = \frac{\rho}{\alpha N(1-\rho)}\mathcal{C}^{(1)}, \qquad \hat{\mathcal{L}}^{(2)} = \frac{\rho}{\alpha N(1-\rho)}\mathcal{C}^{(2)} \tag{14}$$

First, consider $\mathcal{G}^{(1)}$. Note that $\hat{\mathcal{L}}^{(1)} \leq \mathcal{L}^*$, because $\mathcal{G}^{(1)}$ is the set of measurements that $\mathcal{A}_{\text{thb}}[\frac{\rho}{\alpha N(1-\rho)}]$ produces report messages before or at the same time as $\mathtt{Opt}(\sigma_{\mathtt{S}})$ does, and the latency penalty function is non-decreasing in the reporting time. See Fig. 2 (b) for an illustration. Hence,

$$\mathcal{C}^{(1)} = \frac{\alpha N(1-\rho)}{\rho}\hat{\mathcal{L}}^{(1)} \leq \frac{\alpha N(1-\rho)}{\rho}\mathcal{L}^* \tag{15}$$

Therefore, the total cost of $\mathcal{A}_{\text{thb}}$ for $\mathcal{G}^{(1)}$ is bounded by

$$\mathtt{Cost}(\mathcal{A}_{\text{thb}}[\tfrac{\rho}{\alpha N(1-\rho)}, \mathcal{G}^{(1)}]) \tag{16}$$

$$= \rho\mathcal{C}^{(1)} + (1-\rho)\mathcal{L}^{(1)} \leq (1-\rho)(\alpha N+1)\mathcal{L}^* \tag{17}$$

Second, consider $\mathcal{G}^{(2)}$. Let the number of report messages by $\mathcal{A}_{\text{thb}}[\frac{\rho}{\alpha N(1-\rho)}]$ for $\mathcal{G}^{(2)}$ be $\eta^{(2)}$. Note that $\eta^{(2)} \leq N$, because there are at most N systems. Since the communication cost per message is bounded by α, we obtain $\alpha N\mathcal{C}^* \geq \mathcal{C}^{(2)}$. Thus,

$$\mathcal{L}^{(2)} \leq \hat{\mathcal{L}}^{(2)} = \frac{\rho}{\alpha N(1-\rho)}\mathcal{C}^{(2)} \leq \frac{\rho}{1-\rho}\mathcal{C}^* \tag{18}$$

Therefore, the total cost of $\mathcal{A}_{\text{thb}}$ for $\mathcal{G}^{(2)}$ is bounded by

$$\mathtt{Cost}(\mathcal{A}_{\text{thb}}[\tfrac{\rho}{\alpha N(1-\rho)}, \mathcal{G}^{(2)}]) \tag{19}$$

$$= \rho\mathcal{C}^{(2)} + (1-\rho)\mathcal{L}^{(2)} \leq \rho(\alpha N+1)\mathcal{C}^* \tag{20}$$

To sum up, the competitive ratio is obtained by

$$\frac{\mathtt{Cost}(\mathcal{A}_{\text{thb}}[\frac{\rho}{\alpha N(1-\rho)}, \sigma_{\mathtt{S}}])}{\mathtt{Opt}(\sigma_{\mathtt{S}})} \tag{21}$$

$$\leq \frac{(1-\rho)(\alpha N+1)\mathcal{L}^* + \rho(\alpha N+1)\mathcal{C}^*}{\rho\mathcal{C}^* + (1-\rho)\mathcal{L}^*} = \alpha N + 1 \tag{22}$$

■

Remark: Theorem 1 naturally generalizes the 2-competitive deterministic online algorithm for the classical TCP ACK problem in [1] (i.e., the one-system case of 1-sDIA). We will show in Sec. VII that the competitive ratio is indeed optimal in terms of order of magnitude among any deterministic distributed online algorithms, as the lower bound is $\Omega(N)$.

Theorem 2: For K-DIA, let threshold $\theta = \frac{K\rho}{\alpha N(1-\rho)}$ for every system $i \in \mathcal{N}$, the competitive ratio of $\mathcal{A}_{\text{thb}}$ is $\mathtt{CR}(\mathcal{A}_{\text{thb}}[\frac{K\rho}{\alpha N(1-\rho)}]) = \frac{\alpha N}{K} + 1$.

Remark: The competitive ratio of $\mathcal{A}_{\text{thb}}$ diminishes with K, because Opt has increased communication cost, while $\mathcal{A}_{\text{thb}}$ produces the same communication cost independent of K.

V. Full Intercommunication

This section considers full intercommunication between any pair of systems. We study the extent of intercommunication can impact on the competitive ratio of online decision making. With the help of collision-free intercommunication, the competitive ratio can be improved from $O(N)$ to $O(\sqrt{N})$, which matches the lower bound $\Omega(\sqrt{N})$ in Sec. VII.

Threshold-based algorithm $\mathcal{A}_{\text{thb}}^{\text{itc},K}$ (Algorithm 2) relies on overhearing between every pair of systems. In an ideal setting, we suppose that there is a priority based collision-free intercommunication mechanism. The local algorithms at systems are invoked sequentially following the priority order, without being invoked simultaneously. Each system i produces a report message when the ratio between the perceived accumulative latency penalty and the accumulative communication cost since the last report message exceeds threshold θ, similar to $\mathcal{A}_{\text{thb}}$. However, each system can overhear the report messages from other systems. It will not produce a report message for those events that have been reported K times. It also adjusts the accumulative latency penalty for the unreported events.

$\mathcal{A}_{\text{thb}}^{\text{itc}.K}$ ensures that each event is reported at least K times, and hence produces a feasible solution for K-DIA.

Let $z_i^j(t)$ be the indicator variable whether system i has reported event j at time t. It is assumed that each event is identifiable in the report message (e.g., by the appearance time, as there is only one event at each time t). Because of the ability of overhearing report messages, the information $(z_{i'}^j(t))_{i' \in \mathcal{N}}$ is shared among all the systems.

Algorithm 2 $\mathcal{A}_{\text{thb}}^{\text{itc}.K}[\theta, t_{\text{now}}, (\sigma_i^j)_{j:\mathtt{t}^j \le t_{\text{now}}}, (z_{i'}^j(t_{\text{now}}))_{i' \in \mathcal{N}}]$

Input: Threshold θ; current time t_{now}; known measurements $(\sigma_i^j)_{j:\mathtt{t}^j \le t_{\text{now}}}$; indicators of reported events $z_{i'}^j(t_{\text{now}})$
Global Init.: $x_k(i,j) \leftarrow 0$ for all $j \in \mathcal{M}, k \leftarrow 0, r_{i,k} \leftarrow 0$

for each j such that $r_{i,k} < \mathtt{t}^j \le t_{\text{now}}$ **do**
  $\text{cnt}_j \leftarrow \sum_{i' \in \mathcal{N}} z_{i'}^j(t_{\text{now}})$ ▷ *Count num. of report messages*
end for
▷ *Compute latency penalty for events reported $< K$ times*
$\mathtt{lat} \leftarrow \sum_{j \in \mathcal{M}: r_{i,k} < \mathtt{t}^j \le t_{\text{now}} \wedge \text{cnt}_j < K} \mathcal{L}(\sigma_i^j, t_{\text{now}})$
▷ *Compute communication cost for events reported $< K$ times*
$\mathtt{com} \leftarrow \mathcal{C}\left(\{\sigma_i^j : j \in \mathcal{M}, r_{i,k} < \mathtt{t}^j \le t_{\text{now}} \wedge \text{cnt}_j < K\}\right)$
if $\frac{\mathtt{lat}}{\mathtt{com}} \ge \theta$ **then** ▷ *Condition on threshold*
  $k \leftarrow k+1$ ▷ *Next report message*
  $x_k(i,j) \leftarrow 1$ for all observed event j
  $z_i^j(t_{\text{now}} + \epsilon) \leftarrow 1$ for all reported event j, and any ϵ
  $r_{i,k} \leftarrow t_{\text{now}}$
end if

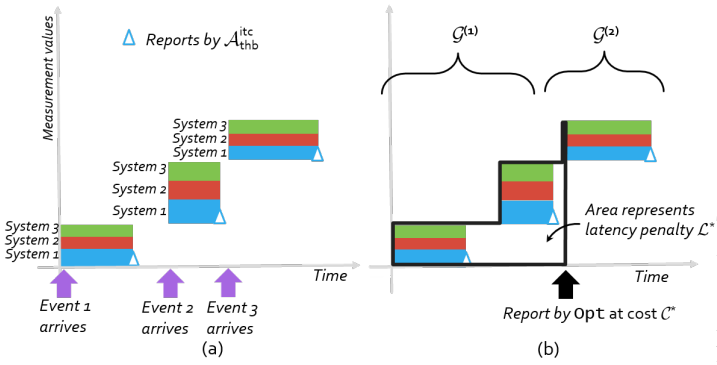

Fig. 3: (a) Considering 1-DIA, $\mathcal{A}_{\text{thb}}^{\text{itc}.1}$ produces one report message per event. (b) Opt produces a report message in the end of each segment $\mathtt{S} \in \mathcal{S}^*$ with communication cost $\mathcal{C}^*$ and latency penalty $\mathcal{L}^*$. The set of measurements is divided into two groups $\mathcal{G}^{(1)}$, $\mathcal{G}^{(2)}$.

Theorem 3: For 1-DIA, let threshold $\theta = \frac{\rho}{(1-\rho)\phi}$ for every system $i \in \mathcal{N}$ where $\phi = \frac{\sqrt{(\alpha-1)^2+4\alpha N}+\alpha-1}{2}$, the competitive ratio $\mathtt{CR}(\mathcal{A}_{\text{thb}}^{\text{itc}.1}[\frac{\rho}{(1-\rho)\phi}]) = \frac{\sqrt{(\alpha-1)^2+4\alpha N}+\alpha+1}{2}$.

Proof: As in Theorem 1, we define $\mathcal{S}^*$, $\mathcal{G}^{(1)}$, $\mathcal{G}^{(2)}$, $\mathcal{C}^*$, $\mathcal{L}^*$, $\mathcal{C}^{(1)}$, $\mathcal{L}^{(1)}$, $\mathcal{C}^{(2)}$, and $\mathcal{L}^{(2)}$. The competitive ratio is obtained with respect to each $\sigma_{\mathtt{S}}$ for $\mathtt{S} \in \mathcal{S}^*$.

First, consider $\mathcal{G}^{(1)}$. Note that because of overhearing, the perceived latency penalty is identical with the actual latency penalty. Hence, $\mathcal{L}^{(1)} \le \mathcal{L}^*$. See Fig. 3 for an illustration.

Note that $\mathcal{A}_{\text{thb}}^{\text{itc}.1}$ produces at most one report message per event. Given $\mathcal{L}^*$, the maximum value of $\mathcal{C}^{(1)}$ is attained when each event is observed by only one system, while other $N-1$ systems incur zero latency. In this case, we obtain

$$\mathcal{C}^{(1)} = \frac{(1-\rho)\phi}{\rho}\mathcal{L}^{(1)} \le \frac{(1-\rho)\phi}{\rho}\mathcal{L}^* \tag{23}$$

Therefore, the total cost of $\mathcal{A}_{\text{thb}}^{\text{itc}.1}$ for $\mathcal{G}^{(1)}$ is bounded by

$$\mathtt{Cost}(\mathcal{A}_{\text{thb}}^{\text{itc}.1}[\tfrac{\rho}{(1-\rho)\phi}, \mathcal{G}^{(1)}]) \le (1-\rho)(\phi+1)\mathcal{L}^* \tag{24}$$

Second, consider $\mathcal{G}^{(2)}$. We bound $\mathcal{L}^{(2)}$ by $\mathcal{C}^{(2)}$. Given $\mathcal{C}^{(2)}$, the maximum value of $\mathcal{L}^{(2)}$ is attained when all N systems observe the same measurement value, while only one report message is generated by one of the systems (because of collision-free intercommunication). Hence,

$$\mathcal{L}^{(2)} \le \frac{N\rho}{(1-\rho)\phi}\mathcal{C}^{(2)} \tag{25}$$

Since there are no simultaneous report messages with collision-free intercommunication, there is at most one report message for $\mathcal{G}^{(2)}$. Hence, $\alpha\mathcal{C}^* \ge \mathcal{C}^{(2)}$. Therefore, the total cost of $\mathcal{A}_{\text{thb}}^{\text{itc}.1}$ for $\mathcal{G}^{(2)}$ is bounded by

$$\mathtt{Cost}(\mathcal{A}_{\text{thb}}^{\text{itc}.1}[\tfrac{\rho}{(1-\rho)\phi}, \mathcal{G}^{(2)}]) \tag{26}$$

$$= \rho\mathcal{C}^{(2)} + (1-\rho)\mathcal{L}^{(2)} \le \alpha\rho\mathcal{C}^* + \alpha\rho\frac{N}{\phi}\mathcal{C}^* \tag{27}$$

To sum up, the competitive ratio is obtained by

$$\frac{\mathtt{Cost}(\mathcal{A}_{\text{thb}}^{\text{itc}.1}[\frac{\rho}{(1-\rho)\phi}, \sigma_{\mathtt{S}}])}{\mathtt{Opt}(\sigma_{\mathtt{S}})} \tag{28}$$

$$\le \frac{(1-\rho)(\phi+1)\mathcal{L}^* + \alpha\rho(\frac{N}{\phi}+1)\mathcal{C}^*}{\rho\mathcal{C}^* + (1-\rho)\mathcal{L}^*} \tag{29}$$

$$= \phi + 1 = \frac{\sqrt{(\alpha-1)^2+4\alpha N}+\alpha+1}{2} \tag{30}$$

The last equality is because that $\phi = \frac{\sqrt{(\alpha-1)^2+4\alpha N}+\alpha-1}{2}$ is the positive root of ϕ in $\phi + 1 = \alpha(\frac{N}{\phi}+1)$. ■

Remark: We note that it is possible to improve the competitive ratio of $\mathcal{A}_{\text{thb}}^{\text{itc}.1}$ by using heterogeneous thresholds. In the proof of Theorem 3, we assume the worst case where the system of maximum communication cost produces a report. One possible approach of improvement is to let the system of minimum communication cost to have a slightly smaller threshold θ_i than other systems, and hence, this system will be triggered to produce a report before others. Let $\theta_i = \theta + \epsilon \cdot \mathcal{C}(\sigma_{i,k})$, where $\mathcal{C}(\sigma_{i,k})$ is communication cost of the next report, and ϵ is an arbitrarily small constant. Hence, the system with a higher communication cost will wait slightly longer. Thus, a report is always produced by the system of minimum communication cost. This is equivalent to having $\alpha = 1$ in Theorem 3, and we set $\theta = \frac{\rho}{(1-\rho)\phi}$, where $\phi = \sqrt{N}$. The competitive ratio of $\mathcal{A}_{\text{thb}}^{\text{itc}.1}$ becomes $\sqrt{N}+1$.

Theorem 4: For K-DIA, let threshold $\theta = \frac{\rho}{(1-\rho)\phi}$ for every system $i \in \mathcal{N}$ where $\phi = \frac{\sqrt{(\alpha-K)^2+4\alpha K N}+\alpha-K}{2K}$, the competitive ratio $\mathtt{CR}(\mathcal{A}_{\text{thb}}^{\text{itc}.K}[\frac{\rho}{(1-\rho)\phi}]) = \frac{\sqrt{(\alpha-K)^2+4\alpha K N}+\alpha+K}{2K}$.

VI. Partial Intercommunication

This section considers non-ideal intercommunication, modelled by a communication graph. A communication graph is defined by $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, where each edge $(i, i') \in \mathcal{E}$ represents that a pair of systems that can engage in full intercommunication (including overhearing) with each other.

We consider a general process that consists of two options:

1) *Overhear-and-withhold*: When a system overhears the report messages from neighboring systems, it will not generate a report message for those events that have been reported K times. It also adjusts the accumulative latency penalty for the unreported events.
2) *Overhear-and-forward*: When a system overhears the report messages from neighboring systems, it will replicate the report messages, combining with its on-going report message to its neighbors. It is assumed that there is collision-free intercommunication with the neighbors.

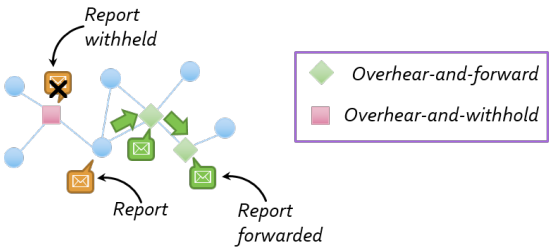

Fig. 4: An illustration for overhear-and-withhold and overhear-and-forward operations.

We partition the systems into the two sets. Let $\mathcal{N}_{\rm w} \subseteq \mathcal{N}$ be the set of overhear-and-withhold nodes, and $\mathcal{N}_{\rm f} = \mathcal{N} \backslash \mathcal{N}_{\rm w}$ be the set of overhear-and-forward nodes. Threshold-based algorithm $\mathcal{A}_{\rm thb}^{{\rm net}.K}$ (Algorithm 3) performs operations according to the role of each node.

Let $\mathtt{x}$ be the maximum number of systems that can produce simultaneous report messages in $\mathcal{G}$. That is,

$$\mathtt{x} \triangleq \max_{\mathcal{X} \subset \mathcal{N}_{\rm w}} \Big| \{ \mathcal{N}_{\rm f} \cup \mathcal{X} \mid i \in \mathcal{X} \Rightarrow \nexists i' \in \mathcal{N}_{\rm f} \cup \mathcal{X} \text{ and } (i, i') \in \mathcal{E} \} \Big|$$

Namely, $\mathtt{x}$ is the number of nodes in $\mathcal{N}_{\rm f}$, and the maximum number of nodes in $\mathcal{N}_{\rm w}$ that are not the neighbors of each other or with a node in $\mathcal{N}_{\rm f}$. Here are a few examples of $\mathtt{x}$:

- (NC) For no intercommunication, we set $\mathcal{N}_{\rm w} = \mathcal{N}$ and $\mathcal{E} = \varnothing$, then $\mathtt{x} = N$.
- (FC) For full intercommunication, we set $\mathcal{N}_{\rm w} = \mathcal{N}$ and $\mathcal{E} = \{(i, i')\}_{i,i' \in \mathcal{N}}$, then $\mathtt{x} = 1$.
- (N1) All nodes carry out overhear-and-withhold: $\mathcal{N}_{\rm w} = \mathcal{N}$. Let the maximum independent set of $\mathcal{G}$ by $\mathtt{MIS}(\mathcal{G})$. Then $\mathtt{x} = |\mathtt{MIS}(\mathcal{G})|$.
- (N2) The nodes that belong to a minimum connected dominating set (CDS) carry out overhear-and-forward, while the rest of nodes carry out overhear-and-withhold. We assume that $\mathcal{G}$ is connected. Let the minimum connected dominating set by $\mathtt{CDS}(\mathcal{G})$. Let $\mathcal{N}_{\rm f} = \mathtt{CDS}(\mathcal{G})$, $\mathcal{N}_{\rm w} = \mathcal{N} \backslash \mathtt{CDS}(\mathcal{G})$. Then $\mathtt{x} = |\mathtt{CDS}(\mathcal{G})|$.

Algorithm 3 $\mathcal{A}_{\rm thb}^{{\rm net}.K}[\theta, t_{\rm now}, (\sigma_i^j)_{j:\mathtt{t}^j \le t_{\rm now}}, (z_{i'}^j(t_{\rm now}))_{i' \in \mathcal{N}:(i,i') \in \mathcal{E}}]$

Input: Threshold θ; current time $t_{\rm now}$; known measurements $(\sigma_i^j)_{j:\mathtt{t}^j \le t_{\rm now}}$; indicators of reported events $z_{i'}^j(t_{\rm now})$ from neighbors.

Global Init.: $x_k(i,j) \leftarrow 0$ for all $j \in \mathcal{M}, k \leftarrow 0, r_{i,k} \leftarrow 0$

▷ *Count num. of reports from neighbors*
for each j such that $r_{i,k} < \mathtt{t}^j \le t_{\rm now}$ **do**
  $\mathrm{cnt}_j \leftarrow \sum_{i':(i,i') \in \mathcal{E}} z_{i'}^j(t_{\rm now})$
end for
▷ *Compute latency penalty for events reported $< K$ times*
$\mathtt{lat} \leftarrow \sum_{j \in \mathcal{M}: r_{i,k} < \mathtt{t}^j \le t_{\rm now} \wedge \mathrm{cnt}_j < K} \mathcal{L}(\sigma_i^j, t_{\rm now})$
▷ *Compute communication cost for events reported $< K$ times*
$\mathtt{com} \leftarrow \mathcal{C}\Big(\{\sigma_i^j : j \in \mathcal{M}, r_{i,k} < \mathtt{t}^j \le t_{\rm now} \wedge \mathrm{cnt}_j < K\}\Big)$
if $\frac{\mathtt{lat}}{\mathtt{com}} \ge \theta$ **then** ▷ *Condition on threshold*
  if $i \in \mathcal{N}_{\rm f}$ **then** ▷ *Overhear-and-forward*
    $k \leftarrow k + 1$ ▷ *Next report message*
    $x_k(i,j) \leftarrow 1$ for all observed j or if $z_{i'}^j(t_{\rm now}) = 1$
    $z_i^j(t_{\rm now} + \epsilon) \leftarrow 1$ for all reported j or if $z_{i'}^j(t_{\rm now}) = 1$
    $r_{i,k} \leftarrow t_{\rm now}$
  else if $i \in \mathcal{N}_{\rm w}$ **then** ▷ *Overhear-and-withhold*
    $k \leftarrow k + 1$ ▷ *Next report message*
    $x_k(i,j) \leftarrow 1$ for all observed event j
    $z_i^j(t_{\rm now} + \epsilon) \leftarrow 1$ for all reported event j, and any ϵ
    $r_{i,k} \leftarrow t_{\rm now}$
  end if
end if

Theorem 5: Given a communication graph $\mathcal{G}$, with $\mathcal{N}_{\rm w}$ and $\mathcal{N}_{\rm f}$, Setting threshold $\theta = \frac{\rho}{(1-\rho)\phi}$ for every system $i \in \mathcal{N}$, where

$$\phi = \frac{\sqrt{(\alpha \mathtt{x} - K)^2 + 4\alpha K N} + \alpha \mathtt{x} - K}{2K} \tag{31}$$

the competitive ratio of $\mathcal{A}_{\rm thb}^{{\rm net}.K}$ is

$$\mathtt{CR}(\mathcal{A}_{\rm thb}^{{\rm net}.K}[\tfrac{\rho}{(1-\rho)\phi}]) = \frac{\sqrt{(\alpha \mathtt{x} - K)^2 + 4\alpha K N} + \alpha \mathtt{x} + K}{2K} \tag{32}$$

Remark: For no intercommunication ($\mathtt{x} = N$), it follows that the competitive ratio becomes $\mathtt{CR}(\mathcal{A}_{\rm thb}^{{\rm net}.K}[\frac{\rho}{(1-\rho)\phi}]) = \frac{\alpha N}{K} + 1$ $(= \mathtt{CR}(\mathcal{A}_{\rm thb}[\frac{K\rho}{\alpha N(1-\rho)}]))$. For full intercommunication ($\mathtt{x} = 1$), it follows that the competitive ratio becomes $\mathtt{CR}(\mathcal{A}_{\rm thb}^{{\rm net}.K}[\frac{\rho}{(1-\rho)\phi}]) = \frac{\sqrt{(\alpha-K)^2 + 4\alpha K N} + \alpha + K}{2K}$ $(= \mathtt{CR}(\mathcal{A}_{\rm thb}^{{\rm itc}.K}[\frac{\rho}{(1-\rho)\phi}]))$. Thus, $\mathcal{A}_{\rm thb}^{{\rm net}.K}[\frac{\rho}{(1-\rho)\phi}]$ provides a continuous transition between the two cases.

VII. Lower Bounds of Competitive Ratios

This section shows that the threshold-based algorithms ($\mathcal{A}_{\rm thb}$ and $\mathcal{A}_{\rm thb}^{{\rm itc}.K}$) achieve the best competitive ratios in terms of order of magnitude among all deterministic distributed online algorithms in the respective settings. It suffices to consider only 1-DIA for establishing the lower bounds of competitive ratios.

In practice, intercommunication consumes at least one transmission but still may require an extra transmission for reporting. We assume that *any intercommunication among systems is as costly as transmitting a report message*.

Theorem 6: Consider 1-DIA. Suppose $\rho = \frac{1}{2}$ and there is one system having a communication cost per report as 1, while other $(N-1)$ systems having a communication cost per report as $\alpha \geq 1$. There exists no deterministic distributed online algorithm with no intercommunication to achieve a competitive ratio lower than $\alpha(N-1) + \frac{3}{2}$.

Theorem 7: Consider 1-DIA and $\rho = \frac{1}{2}$. There exists no deterministic distributed online algorithm with collision-free intercommunication to achieve a competitive ratio lower than $\frac{\sqrt{N}}{4}$.

Remark: Theorems 6-7 show that the competitive ratios for threshold-based algorithms in Theorem 1 and Theorem 3 (i.e., $O(N)$ and $O(\sqrt{N})$ respectively) are optimal in order of magnitude.

VIII. Simulation Studies

We provided analysis on the competitive ratios of our algorithms in the previous sections, which are the worst-case guarantees. In this section, we evaluate the empirical *average-case* ratios by simulations. We compare the average-case ratios with our theoretical competitive ratios, and observe that our algorithms perform relatively well in several scenarios which are far below the theoretical worst-case values.

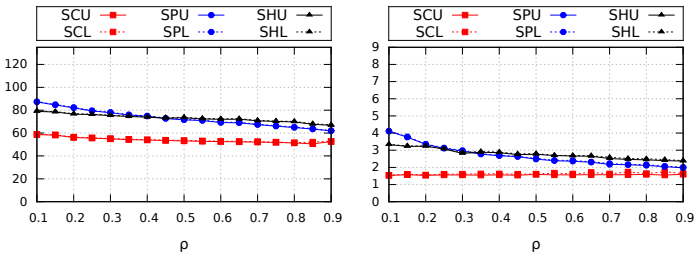

(a) No intercommunication (b) Full intercommunication

Fig. 6: The x-axis corresponds to the different values of ρ, whereas the y-axis is the empirical ratio.

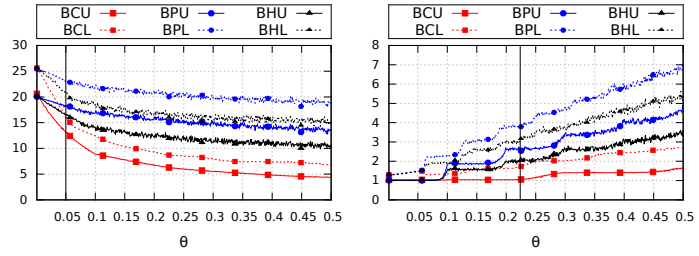

(a) No intercommunication (b) Full intercommunication

Fig. 7: The x-axis corresponds to the different values of θ, whereas the y-axis is the empirical ratio. The solid vertical lines correspond to the values of θ derived by our theorems.

A. Simulation Set-up

The simulations were performed up to 300 systems, and each has 50 runs. For the partial intercommunication, we generate a random connected unit disk graphs with a fixed average degree (≈ 18) for each run. Fig. 9a shows an example of a random network of 100 systems in the simulation.

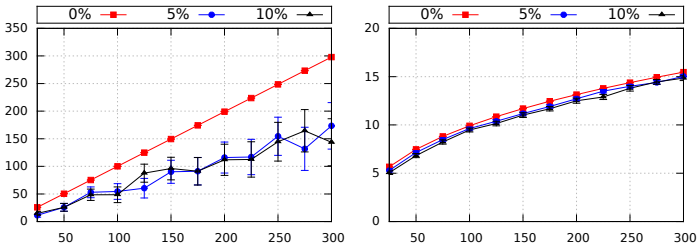

(a) No intercommunication (b) Full intercommunication

Fig. 8: Simulations of randomly perturbed worst-case inputs, where the measurements are perturbed by a random percent. The x-axis corresponds to the different values of N, whereas the y-axis is the empirical ratio.

The measurement values (w_j^i) are uniformly distributed generated by the following two ranges:

1) *Big events* (B): $w_j^i \in [0, 1]$ for all i, j.
2) *Small events* (S): $w_j^i \in [0, \frac{1}{N}]$ for all i, j.

The appearance times $(\mathfrak{t}^j)$ are generated by the following three random arrival models:

1) *Constant-interval arrivals* (C): The inter-arrival time is a constant.
2) *Poisson arrivals* (P): The inter-arrival time follows an exponential distribution with mean of 20 time slots.
3) *Heavy-tailed arrivals* (H): The inter-arrival time follow Weibull distribution with shape parameter 0.5, and mean of 10 time slots.

The latency penalty function is chosen to be linear, while the following two communication cost functions are considered:

1) *Unity cost* (U): The communication cost is always 1, $\mathcal{C}(\sigma_{i,k}) = 1$.
2) *Logarithmic cost* (L): The communication cost is a logarithm function of local total measurement value, $\mathcal{C}(\sigma_{i,k}) = \log\left(2 + \sum_{j \in \mathcal{M}: \sigma_i^j \in \sigma_{i,k}} w_i^j\right)$.

In the following, a simulation scenario is represented by the acronyms. For instance, BHL represents the scenario with big events, heavy-tailed arrivals and logarithmic cost.

B. Observations and Discussion

We present the simulation results for each scenario, and discuss the observations.

1) **No Intercommunication**: Figs. 5a-5b show the empirical ratios for different scenarios. Overall, we observe linear trends in the empirical ratios as N increases, which are below the respective theoretical competitive ratios. Note that constant-interval arrivals (C) give consistently better empirical ratios, followed by Poisson arrivals (P) and heavy-tailed arrivals (H). However, when using big events (Fig. 5b), the difference in empirical ratio is much smaller. We also observe that logarithmic cost (L) results in higher empirical ratios than unity cost which is justified by the additional factor α in Theorem 1.

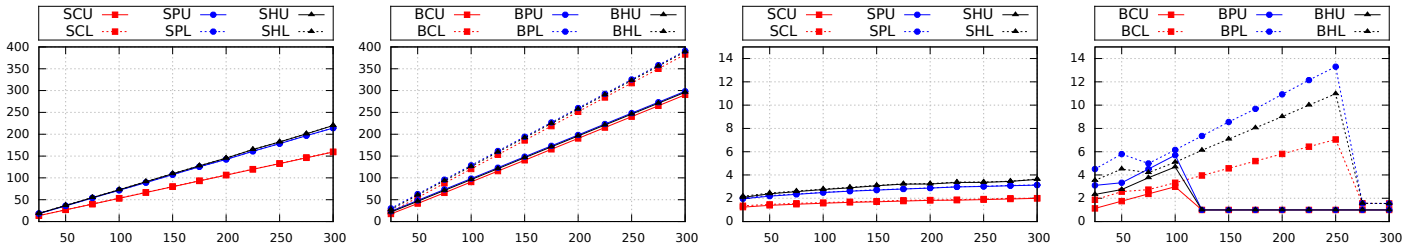

(a) No intercommunication (b) No intercommunication (c) Full intercommunication (d) Full intercommunication

Fig. 5: Different event arrival models for the settings of no intercommunication and full intercommunication. The y-axis is the empirical ratio, whereas the x-axis is the number of systems N.

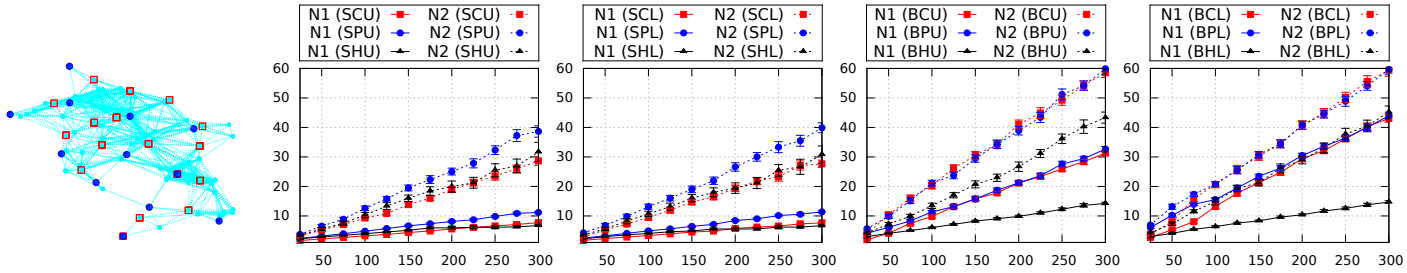

(a) A random UDG (b) Unity cost (c) Log cost (d) Unity cost (e) Log cost

Fig. 9: (a) An example of a random Unit Disk Graph (UDG) of 100 nodes that is generated randomly in the simulation. The red squares represent a CDS, while the blue circles belong to MIS. (b-e) The figures show different scenarios for the setting of partial intercommunication. The y-axis is the empirical ratio, whereas the x-axis is the number of systems N.

We present a sensitivity analysis for various values of ρ and θ. We set the number of systems as 100. Fig. 6a shows that higher value of ρ achieves the better (i.e., lower) empirical ratio. Fig. 7a illustrates the effect of using different values of θ rather than the one derived by our theorems, which is shown by a vertical line. The higher value θ achieves the better empirical ratio.

Fig.8a shows the empirical ratios of randomly perturbed worst-case input from Theorem 6, where the measurements are perturbed by a random percent. Slightly perturbed inputs significantly reduce the empirical ratio. Fig. 10a shows the empirical ratios for the K-report problem when $K = 1, 2, 3$ for scenario SPU. A larger K yields a lower empirical ratio, as shown in Theorem 2.

2) **Full Intercommunication**: Figs. 5c-5d show the empirical ratio for different scenarios. Overall, we observe a sub-linear trend in the empirical ratio as N increases, which are below the theoretical competitive ratios. For small events, the empirical ratios increase smoothly. However, for big events, the empirical ratios drop to 1, as N reaches a certain value. The intuition for such a drop is that when N is large, the total latency penalty exceeds the communication cost for almost all events j at every time slot, therefore an optimal offline solution generates a report message at every time slot, which coincides with the threshold-based algorithm.

Fig. 6b shows that different values of ρ has small effect on the empirical ratio for small events. Fig. 7b illustrates the effect of using different values of θ. Increasing θ beyond that is derived by our theorem increases the empirical ratio. Fig.8b shows the empirical ratios of randomly perturbed worst-case input from Theorem 7.

Fig. 10b shows the empirical ratios for the K-report problem when $K = 1, 2, 3$ for scenario SPU. A larger K yields a lower empirical ratio.

3) **Partial Intercommunication**: Figs. 9b-9e show the empirical ratios for partial intercommunication using settings of N1 and N2 in Sec. VI. We observe that N1 obtains smaller empirical ratios than N2 for all scenarios. We used approximation algorithms to obtain MIS and CDS. In most scenarios we observe that CDS is slightly larger than MIS (see Fig. 9a for an example), which creates a larger empirical ratio for N2. We also observe that small events give smaller empirical ratios than large events in all scenarios. Moreover, logarithmic cost gives larger empirical ratios as expected.

IX. Practical Sensor Network Testbed

As a concrete demonstration for our distributed online algorithms, we implemented our threshold-based algorithms in a practical sensor network testbed for applications such as intrusion detection and air-conditioning control. For each application, we deploy a network of wireless sensor nodes equipped with appropriate sensors for detecting the relevant signals from the environment, which then transmit report messages to a base station.

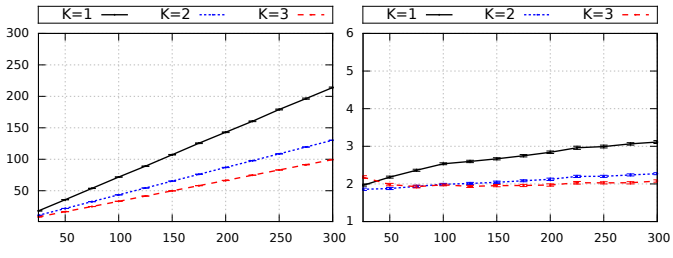

(a) No comm. (SPU) (b) Full comm. (SPU)

Fig. 10: The effect of different values of K on the empirical ratio for the settings of no intercommunication and full intercommunication for scenario SPU.

A. Hardware Set-up

The sensor network testbed is deployed in an office building (see Fig. 11). The testbed consists of:

1) A network of wireless sensor nodes programmed with our threshold-based algorithms. We use several types of sensors to monitor various events depending on the applications. For example, the sound sensor is an XBee-enabled Arduino board equipped with a sound detector while the remaining sensors are TinyOS based. Each sensor can be set with broadcast mode, which can overhear other sensors' transmissions.
2) Several TinyOS-based base stations that are also connected to a laptop computer and record data from received from TinyOS based sensor nodes.
3) Several Arduino boards equipped with buzzers and lights that can act as random fluctuating sound and light sources, and generate measurement for the sound sensors and light sensors.
4) A room heater controlled by Arduino that acts as random fluctuating heat source to generate measurements for temperature sensors.
5) A hardware module that monitors the energy consumptions of wireless sensors.

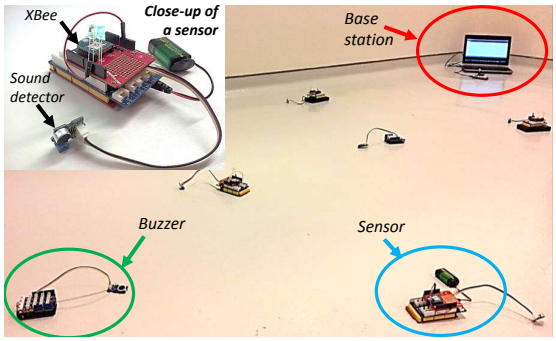

Fig. 11: A snapshot of the sensor network testbed.

To implement collision-free intercommunication, we use a simple priority based mechanism. Each sensor node is assigned priority order in advanced, which determines for how long the sensor must wait (in multiple of 10ms) before transmitting data over the radio channel. The configurations for each sensor are remotely controllable from the base station.

B. Empirical Results and Discussion

We present the empirical evaluation of the competitive ratios of our threshold-based algorithms. We conducted the experiments for 1-report problem with (1) no intercommunication, where the wireless broadcast mode is off, and (2) full intercommunication, where the wireless broadcast mode is on.

Each wireless sensor is programmed to take measurements at appropriate intervals. Whenever the accumulative latency as a function of measured signal exceeds the specified threshold in our algorithms, the wireless sensor will transmit a notification to the base station via ZigBee wireless link. In case the broadcast mode is turned on and the accumulative latency exceeds the threshold for multiple sensors simultaneously, only the sensor with highest priority will transmit the report.

We next describe the empirical observations. For each type of sensor, we depict a sample trace of the measured signals and plot the results in terms of competitive ratios in Fig. 12. To compute the competitive ratios, we divide the total cost by the respective optimal offline cost.

1) Sensor Data: Sound sensors can detect the intensity of loudness of environmental sound and generate an output in the range (0-1023) accordingly. Light sensors can detect the intensity of brightness in the range (0-4015). Motion sensors can generate binary states (i.e., 1 whenever motion is detected, 0 otherwise). Figs. 12a, 12c and 12e show the sample traces of the sound, light, and motion sensors respectively. We observe that these sensors have rather instantaneous response time, which makes them suitable for intrusion detection applications.

Figs. 12b and 12d show the signal traces for temperature and CO_2 sensors. Unlike the previous sensors, the response time of both temperature and CO_2 sensors is not slower due to the slowly varying concentration of CO_2 level. We observe that in the absence of external CO_2 sources (e.g., human), the values measured by the sensors are around 500 ppm, which is close to the air concentration of CO_2 under normal conditions.

2) Sensor Energy Consumption: One of the design goals in wireless sensor networks is prolong the battery lifetime. The abstract communication cost in our model can naturally capture the energy consumption in communication operations. To measure the energy consumption, we conducted a sensor energy monitoring experiment, using a custom-built hardware module that monitors the energy consumption of the wireless sensors used in our experiments. The Arduino-based hardware module measures the energy consumption of various operations performed by wireless sensors such as communication, sensing, and processing. The energy consumption due to communication was measured to be 10 micro-Joules per bit, whereas the energy consumed by sensing and processing operations was observed to be negligible. In our experiments, the radio packet size was fixed at 10 bytes, which means that the wireless sensors consume about 0.8 milli-Joules energy per radio transmission.

3) Empirical Results: Figs. 12a-12e show the observed ratios achieved by our threshold-based online algorithms in the sensor network testbed. There are several observations:

- The observed ratio in each case is considerably below the corresponding theoretical competitive ratio.

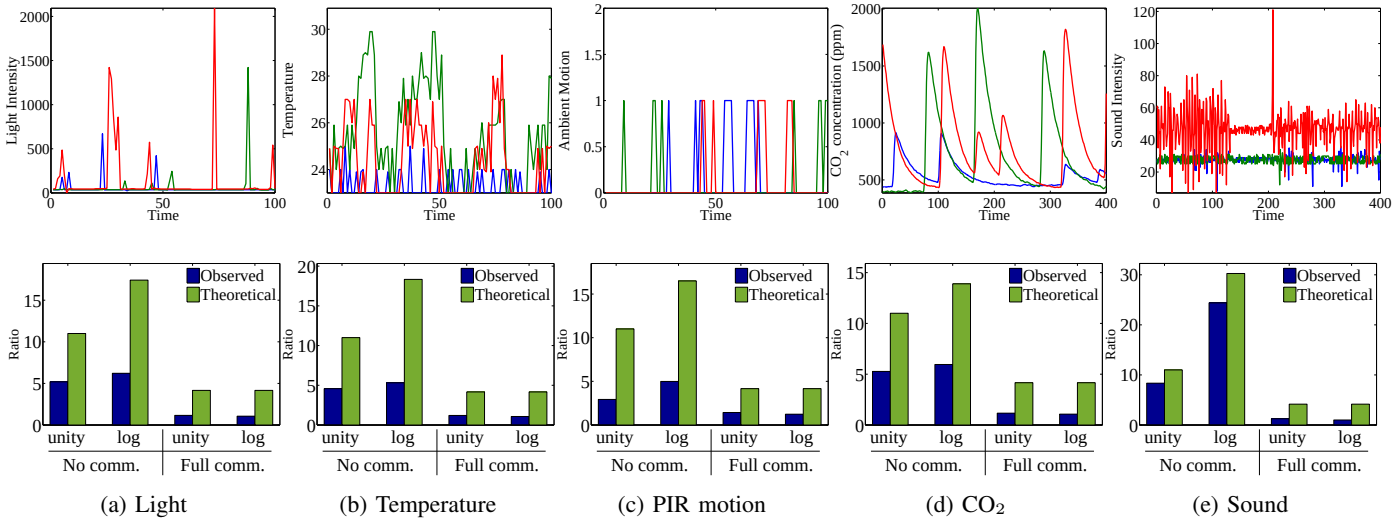

(a) Light (b) Temperature (c) PIR motion (d) CO_2 (e) Sound

Fig. 12: Testbed results

- The observed ratio achieved in the setting of full intercommunication is better than the setting of no intercommunication, because in the former case the number of transmissions by the online algorithm is near optimal as all sensors can overhear each other radio transmissions. In the latter case, on the other hand, the algorithm tends to transmit more reports because the sensors cannot overhear each other.
- Unity cost results in a lower observed ratio than logarithmic cost which is verified by our simulation results.

X. Conclusion

In this paper, we shed light on the distributed online optimization of cost efficiency and effectiveness of information aggregation. In particular, we present competitive online algorithms and theoretical analysis for the trade-off optimization of general communication cost and latency penalty with distributed systems in various settings of intercommunication. All the proofs can be found in the technical report [3].

In a companion paper [16], we also study randomized online algorithms to solve the distributed online optimization, which relies on a different formulation using LP approach.

Distributed online algorithms are a burgeoning area, which is vital to many applications in networking and distributed systems. Our study concerns a general problem, paving the way of tackling a wide range of distributed online problems.

References

[1] D. R. Dooly, S. A. Goldman, and S. D. Scott, "On-line analysis of the TCP acknowledgment delay problem," *J. ACM*, vol. 48, no. 2, pp. 243–273, Mar. 2001.

[2] A. Borodin and R. El-Yaniv, *Online computation and competitive analysis.* Cambridge University Press Cambridge, 1998.

[3] C.-K. Chau, M. Khonji, and M. Aftab, "Online algorithms for information aggregation from distributed and correlated sources," Tech. Rep., 2016, http://arxiv.org/abs/1601.03147.

[4] C.-K. Chau, G. Zhang, and M. Chen, "Cost minimizing online algorithms for energy storage management with worst-case guarantee," *to appear in IEEE Transactions on Smart Grid*, 2016, http://arxiv.org/abs/1511.07559.

[5] L. Lu, J. Tu, C.-K. Chau, M. Chen, and X. Lin, "Online energy generation scheduling for microgrids with intermittent energy sources and co-generation," in *ACM SIGMETRICS*, 2013.

[6] M. Aftab, C.-K. Chau, and P. Armstrong, "Smart air-conditioning control by wireless sensors: An online optimization approach," in *ACM Intl. Conf. on Future Energy Systems (e-Energy)*, 2013.

[7] A. R. Karlin, C. Kenyon, and D. Randall, "Dynamic TCP acknowledgement and other stories about e/(e-1)," in *ACM Symposium on Theory of Computing (STOC)*, 2001.

[8] N. Buchbinder, K. Jain, and J. S. Naor, "Online primal-dual algorithms for maximizing ad-auctions revenue," in *European Conf. on Algorithms (ESA)*, 2007.

[9] Z. Lotker, B. Patt-Shamir, and D. Rawitz, "Rent, lease or buy: Randomized algorithms for multislope ski rental," in *Symposium on the Theoretical Aspects of Computer Science (STACS)*, 2008.

[10] A. Khanafer, M. S. Kodialam, and K. P. N. Puttaswamy, "The constrained ski-rental problem and its application to online cloud cost optimization." in *IEEE INFOCOM*, 2013.

[11] E. Fasolo, M. Rossi, J. Widmer, and M. Zorzi, "In-network aggregation techniques for wireless sensor networks: a survey," *IEEE Wireless Communications*, vol. 14, no. 2, pp. 70–87, April 2007.

[12] C. F. Brito, E. Koutsoupias, and S. Vaya, "Competitive analysis of organization networks or multicast acknowledgment: How much to wait?" *Algorithmica*, vol. 64, no. 4, pp. 584–605, 2012.

[13] S. Khanna, J. S. Naor, and D. Raz, "Control message aggregation in group communication protocols," in *Automata, Languages and Programming*, ser. Lecture Notes in Computer Science. Springer-Verlag, 2002, vol. 2380, pp. 135–146.

[14] P. Korteweg, A. Marchetti-Spaccamela, L. Stougie, and A. Vitaletti, "Data aggregation in sensor networks: Balancing communication and delay costs," in *Structural Information and Communication Complexity*, ser. Lecture Notes in Computer Science. Springer-Verlag, 2007, vol. 4474, pp. 139–150.

[15] L. Becchetti, P. Korteweg, A. Marchetti-Spaccamela, M. Skutella, L. Stougie, and A. Vitaletti, "Latency constrained aggregation in sensor networks," in *Algorithms ESA*, ser. Lecture Notes in Computer Science. Springer-Verlag, 2006, vol. 4168, pp. 88–99.

[16] M. Khonji and C.-K. Chau, "Randomized online algorithms for in-situ sensing information aggregation," Masdar Institute, Tech. Rep., 2016.

APPENDIX

A. Proofs

Lemma 1: (Subadditivity) For a set of segments $\mathcal{S}$,

$$\texttt{Cost}(\mathcal{A}_{\text{thb}}[\theta,\sigma]) \leq \sum_{\texttt{s}\in\mathcal{S}} \texttt{Cost}(\mathcal{A}_{\text{thb}}[\theta,\sigma_{\texttt{s}}]) \tag{33}$$

Subadditivity property (Eqn. (33)) also holds for threshold-based algorithms $\mathcal{A}_{\text{thb}}^{\text{itc}.K}$ and $\mathcal{A}_{\text{thb}}^{\text{net}.K}$.

Proof: It suffices to consider one system. Without loss of generality, we consider a pair of consecutive segments $\texttt{S}_1, \texttt{S}_2 \in \mathcal{S}$. Let $\texttt{Cost}(\mathcal{A}_{\text{thb}}[\theta,\sigma_{\texttt{S}_1\cup\texttt{S}_2}]) = \rho\mathcal{C}_{\texttt{S}_1\cup\texttt{S}_2} + (1-\rho)\mathcal{L}_{\texttt{S}_1\cup\texttt{S}_2}$, $\texttt{Cost}(\mathcal{A}_{\text{thb}}[\theta,\sigma_{\texttt{S}_1}]) = \rho\mathcal{C}_{\texttt{S}_1} + (1-\rho)\mathcal{L}_{\texttt{S}_1}$ and $\texttt{Cost}(\mathcal{A}_{\text{thb}}[\theta,\sigma_{\texttt{S}_2}]) = \rho\mathcal{C}_{\texttt{S}_2} + (1-\rho)\mathcal{L}_{\texttt{S}_2}$.

Note that the communication cost is sub-additive. Thus,

$$\mathcal{C}_{\texttt{S}_1\cup\texttt{S}_2} \leq \mathcal{C}_{\texttt{S}_1} + \mathcal{C}_{\texttt{S}_2} \tag{34}$$

By threshold-based algorithm, one obtains

$$\frac{\mathcal{L}_{\texttt{S}_1\cup\texttt{S}_2}}{\mathcal{C}_{\texttt{S}_1\cup\texttt{S}_2}} = \frac{\mathcal{L}_{\texttt{S}_1}}{\mathcal{C}_{\texttt{S}_1}} = \frac{\mathcal{L}_{\texttt{S}_2}}{\mathcal{C}_{\texttt{S}_2}} = \theta \tag{35}$$

Hence,

$$\mathcal{L}_{\texttt{S}_1\cup\texttt{S}_2} = \theta\mathcal{C}_{\texttt{S}_1\cup\texttt{S}_2} \leq \theta(\mathcal{C}_{\texttt{S}_1} + \mathcal{C}_{\texttt{S}_2}) = \mathcal{L}_{\texttt{S}_1} + \mathcal{L}_{\texttt{S}_2} \tag{36}$$

Therefore,

$$\texttt{Cost}(\mathcal{A}_{\text{thb}}[\theta,\sigma_{\texttt{S}_1\cup\texttt{S}_2}]) \leq \texttt{Cost}(\mathcal{A}_{\text{thb}}(\theta,\sigma_{\texttt{S}_1})) + \texttt{Cost}(\mathcal{A}_{\text{thb}}[\theta,\sigma_{\texttt{S}_2}])$$

Similarly, the preceding argument can be applied to threshold-based algorithms $\mathcal{A}_{\text{thb}}^{\text{itc}.K}$ and $\mathcal{A}_{\text{thb}}^{\text{net}.K}$. ∎

Theorem 2: For K-DIA, let threshold $\theta = \frac{K\rho}{\alpha N(1-\rho)}$ for every system $i \in \mathcal{N}$, the competitive ratio of $\mathcal{A}_{\text{thb}}$ is $\texttt{CR}(\mathcal{A}_{\text{thb}}[\frac{K\rho}{\alpha N(1-\rho)}]) = \frac{\alpha N}{K} + 1$.

Proof: Let $\texttt{Opt}_K(\sigma)$ be the K-report optimal offline solution, and its total cost be $\rho\mathcal{C}_K^* + (1-\rho)\mathcal{L}_K^*$.

Next, consider a different communication cost function, defined by $\tilde{\mathcal{C}}(\sigma_{i,k}) \triangleq K \cdot \mathcal{C}(\sigma_{i,k})$. Let $\widetilde{\texttt{Opt}}(\sigma)$ be the 1-report optimal offline solution using communication cost function $\tilde{\mathcal{C}}(\sigma_{i,k})$, and its total cost be $\rho\tilde{\mathcal{C}}^* + (1-\rho)\tilde{\mathcal{L}}^*$.

$\widetilde{\texttt{Opt}}(\sigma)$ can be compared with $\texttt{Opt}_1(\sigma)$, such that each report message is replicated K times from the same system that has the minimal communication cost. On the other hand, $\texttt{Opt}_K(\sigma)$ requires K report messages from different systems. Since K messages can be transmitted at the system with the least communication cost for $\widetilde{\texttt{Opt}}(\sigma)$, we obtain

$$\rho\mathcal{C}_K^* + (1-\rho)\mathcal{L}_K^* \geq \rho\tilde{\mathcal{C}}^* + (1-\rho)\tilde{\mathcal{L}}^* \tag{37}$$

Therefore, the competitive ratio of $\mathcal{A}_{\text{thb}}$ is bounded by

$$\frac{\texttt{Cost}(\mathcal{A}_{\text{thb}}[\frac{K\rho}{\alpha N(1-\rho)},\sigma])}{\texttt{Opt}_K(\sigma)} \leq \frac{\texttt{Cost}(\mathcal{A}_{\text{thb}}[\frac{K\rho}{\alpha N(1-\rho)},\sigma])}{\widetilde{\texttt{Opt}}(\sigma)} \tag{38}$$

In the following, we consider $\widetilde{\texttt{Opt}}(\sigma)$. Note that $\mathcal{A}_{\text{thb}}$ behaves similarly for K-DIA. As in Theorem 1, we define $\mathcal{S}^*$, $\mathcal{G}^{(1)}$, $\mathcal{G}^{(2)}$, $\mathcal{C}^{(1)}$, $\mathcal{L}^{(1)}$, $\mathcal{C}^{(2)}$, and $\mathcal{L}^{(2)}$. Next, the competitive ratio is obtained with respect to each $\sigma_{\texttt{S}}$ for $\texttt{S} \in \mathcal{S}^*$. By a slight abuse of notation, let the total cost of $\widetilde{\texttt{Opt}}(\sigma_{\texttt{S}})$ be $\rho\tilde{\mathcal{C}}^* + (1-\rho)\tilde{\mathcal{L}}^*$.

For $\mathcal{G}^{(1)}$, as in Theorem 1, we obtain

$$\texttt{Cost}(\mathcal{A}_{\text{thb}}[\tfrac{K\rho}{\alpha N(1-\rho)},\mathcal{G}^{(1)}]) \leq (1-\rho)(\frac{\alpha N}{K}+1)\tilde{\mathcal{L}}^* \tag{39}$$

However, for $\mathcal{G}^{(2)}$, we obtain $\frac{\alpha N}{K}\tilde{\mathcal{C}}^* \geq \mathcal{C}^{(2)}$. Therefore,

$$\texttt{Cost}(\mathcal{A}_{\text{thb}}[\tfrac{K\rho}{\alpha N(1-\rho)},\mathcal{G}^{(2)}]) \leq \rho(\frac{\alpha N}{K}+1)\tilde{\mathcal{C}}^* \tag{40}$$

Finally, it follows that

$$\frac{\texttt{Cost}(\mathcal{A}_{\text{thb}}[\frac{K\rho}{\alpha N(1-\rho)},\sigma_{\texttt{S}}])}{\widetilde{\texttt{Opt}}(\sigma_{\texttt{S}})} \tag{41}$$

$$\leq \frac{(1-\rho)(\frac{\alpha N}{K}+1)\tilde{\mathcal{L}}^* + \rho(\frac{\alpha N}{K}+1)\tilde{\mathcal{C}}^*}{\rho\tilde{\mathcal{C}}^* + (1-\rho)\tilde{\mathcal{L}}^*} = \frac{\alpha N}{K}+1 \tag{42}$$

∎

Theorem 4: For K-DIA, let threshold $\theta = \frac{\rho}{(1-\rho)\phi}$ for every system $i \in \mathcal{N}$ where $\phi = \frac{\sqrt{(\alpha-K)^2+4\alpha KN}+\alpha-K}{2K}$, the competitive ratio $\texttt{CR}(\mathcal{A}_{\text{thb}}^{\text{itc}.K}[\frac{\rho}{(1-\rho)\phi}]) = \frac{\sqrt{(\alpha-K)^2+4\alpha KN}+\alpha+K}{2K}$.

Proof: As in Theorem 2, we define $\tilde{\mathcal{C}}(\sigma_{i,k}) \triangleq K \cdot \mathcal{C}(\sigma_{i,k})$. Let $\widetilde{\texttt{Opt}}(\sigma)$ be the 1-report optimal offline solution using communication cost function $\tilde{\mathcal{C}}(\sigma_{i,k})$. As in Theorem 1, we also define $\mathcal{S}^*$ with respect to $\widetilde{\texttt{Opt}}(\sigma)$, $\tilde{\mathcal{C}}^*$, $\tilde{\mathcal{L}}^*$, $\mathcal{G}^{(1)}$, $\mathcal{G}^{(2)}$, $\mathcal{C}^{(1)}$, $\mathcal{L}^{(1)}$ and $\mathcal{L}^{(2)}$. Next, the competitive ratio is obtained with respect to each $\sigma_{\texttt{S}}$ for $\texttt{S} \in \mathcal{S}^*$.

As in Theorem 3, the total cost of $\mathcal{A}_{\text{thb}}^{\text{itc}.K}$ for $\mathcal{G}^{(1)}$ is

$$\texttt{Cost}(\mathcal{A}_{\text{thb}}^{\text{itc}.K}[\tfrac{\rho}{(1-\rho)\phi},\mathcal{G}^{(1)}]) \leq (1-\rho)(\phi+1)\tilde{\mathcal{L}}^* \tag{43}$$

Let the number of report messages for $\mathcal{G}^{(2)}$ be $\eta^{(2)}$ and the k-th report communication cost be $\mathcal{C}_k^{(2)}$. Since $\widetilde{\texttt{Opt}}(\sigma)$ generates only one report message at communication cost $\tilde{\mathcal{C}}^*$, which is K-time of the per message communication cost of $\mathcal{A}_{\text{thb}}^{\text{itc}.K}$, it follows that $\max_{k\in\{1,\ldots,\eta^{(2)}\}} \mathcal{C}_k^{(2)} \leq \frac{\alpha\tilde{\mathcal{C}}^*}{K}$, and

$$\frac{(1-\rho)\phi}{\rho}\mathcal{L}^{(2)} = \sum_{k=1}^{\eta^{(2)}} \mathcal{C}_k^{(2)} \leq \alpha N\frac{\tilde{\mathcal{C}}^*}{K} \tag{44}$$

Therefore, the total cost of $\mathcal{A}_{\text{thb}}^{\text{itc}.K}$ for $\mathcal{G}^{(2)}$ is bounded by

$$\texttt{Cost}(\mathcal{A}_{\text{thb}}^{\text{itc}.K}[\tfrac{\rho}{(1-\rho)\phi},\mathcal{G}^{(2)}]) \leq \frac{\alpha\rho}{K}(\frac{N}{\phi}+1)\tilde{\mathcal{C}}^* \tag{45}$$

To sum up, the competitive ratio is obtained as follows.

$$\frac{\texttt{Cost}(\mathcal{A}_{\text{thb}}^{\text{itc}.K}[\frac{\rho}{(1-\rho)\phi},\sigma_{\texttt{S}}])}{\texttt{Opt}(\sigma_{\texttt{S}})} \tag{46}$$

$$\leq \frac{(1-\rho)(\phi+1)\mathcal{L}^* + \frac{\alpha\rho}{K}(\frac{N}{\phi}+1)\tilde{\mathcal{C}}^*}{\rho\tilde{\mathcal{C}}^* + (1-\rho)\tilde{\mathcal{L}}^*} \tag{47}$$

$$= \phi+1 = \frac{\sqrt{(\alpha-K)^2+4\alpha KN}+\alpha+K}{2K} \tag{48}$$

The last equality is because that $\phi = \frac{\sqrt{(\alpha-K)^2+4\alpha KN}+\alpha-K}{2K}$ is the positive root of ϕ in $\phi+1 = \frac{\alpha\rho}{K}(\frac{N}{\phi}+1)$. ∎

Theorem 5: Given a communication graph $\mathcal{G}$, with $\mathcal{N}_{\text{w}}$ and $\mathcal{N}_{\text{f}}$, Setting threshold $\theta = \frac{\rho}{(1-\rho)\phi}$ for every system $i \in \mathcal{N}$, where

$$\phi = \frac{\sqrt{(\alpha\text{x}-K)^2+4\alpha KN}+\alpha\text{x}-K}{2K} \tag{49}$$

the competitive ratio of $\mathcal{A}_{\rm thb}^{{\rm net}.K}$ is

$$\texttt{CR}(\mathcal{A}_{\rm thb}^{{\rm net}.K}[\tfrac{\rho}{(1-\rho)\phi}]) = \frac{\sqrt{(\alpha {\rm x}-K)^2+4\alpha KN}+\alpha {\rm x}+K}{2K} \tag{50}$$

Proof: The proof extends the arguments of Theorem 4.

As in Theorem 3, we define $\tilde{\mathcal{C}}(\sigma_{i,k}) \triangleq K \cdot \mathcal{C}(\sigma_{i,k})$. Let $\widetilde{\texttt{Opt}}(\sigma)$ be the 1-report optimal offline solution using communication cost function $\tilde{\mathcal{C}}(\sigma_{i,k})$. As in Theorem 1, we define $\mathcal{S}^*$ with respect to $\widetilde{\texttt{Opt}}(\sigma)$, $\tilde{\mathcal{C}}^*$, $\tilde{\mathcal{L}}^*$, $\mathcal{G}^{(1)}$, $\mathcal{G}^{(2)}$, $\mathcal{C}^{(1)}$, $\mathcal{L}^{(1)}$ and $\mathcal{L}^{(2)}$. Next, the competitive ratio is obtained with respect to each $\sigma_{\tt S}$ for ${\tt S} \in \mathcal{S}^*$.

First, we obtain

$$\mathcal{C}^{(1)} = \frac{(1-\rho)\phi}{\rho}\mathcal{L}^{(1)} \le \frac{(1-\rho)\phi}{\rho}\tilde{\mathcal{L}}^* \tag{51}$$

Therefore, the total cost of $\mathcal{A}_{\rm thb}^{{\rm net}.K}$ for $\mathcal{G}^{(1)}$ is bounded by

$$\texttt{Cost}(\mathcal{A}_{\rm thb}^{{\rm net}.K}[\tfrac{\rho}{(1-\rho)\phi}, \mathcal{G}^{(1)}]) \le (1-\rho)(\phi+1)\tilde{\mathcal{L}}^* \tag{52}$$

Let the number of report messages for $\mathcal{G}^{(2)}$ be $\eta^{(2)}$ and the k-th report communication cost be $\mathcal{C}_k^{(2)}$. As in Theorem 4, it follows that $\max_{k\in\{1,\ldots,\eta^{(2)}\}} \mathcal{C}_k^{(2)} \le \frac{\alpha\tilde{\mathcal{C}}^*}{K}$, and

$$\frac{(1-\rho)\phi}{\rho}\mathcal{L}^{(2)} \le \sum_{k=1}^{\eta^{(2)}} \mathcal{C}_k^{(2)} \le \alpha N \frac{\tilde{\mathcal{C}}^*}{K} \tag{53}$$

On the other hand, since the maximum number of systems that generate simultaneous report messages for $\mathcal{G}^{(2)}$ is x, it follows that $\alpha {\rm x}\tilde{\mathcal{C}}^* \ge \sum_{k=1}^{\eta^{(2)}} \mathcal{C}_k^{(2)}$.

Therefore, the total cost of $\mathcal{A}_{\rm thb}^{{\rm net}.K}$ for $\mathcal{G}^{(2)}$ is bounded by

$$\texttt{Cost}(\mathcal{A}_{\rm thb}^{{\rm net}.K}[\tfrac{\rho}{(1-\rho)\phi}, \mathcal{G}^{(2)}]) \le \alpha\rho(\frac{N}{K\phi}+{\rm x})\tilde{\mathcal{C}}^* \tag{54}$$

Finally, it follows that

$$\frac{\texttt{Cost}(\mathcal{A}_{\rm thb}^{{\rm net}.K}[\frac{\rho\phi}{1-\rho}, \sigma_{\tt S}])}{\widetilde{\texttt{Opt}}(\sigma_{\tt S})} \tag{55}$$

$$\le \frac{(1-\rho)(\phi+1)\tilde{\mathcal{L}}^* + \alpha\rho(\frac{N}{K\phi}+{\rm x})\tilde{\mathcal{C}}^*}{\rho\tilde{\mathcal{C}}^* + (1-\rho)\tilde{\mathcal{L}}^*} \tag{56}$$

$$= \phi+1 = \frac{\sqrt{(\alpha {\rm x}-K)^2+4\alpha KN}+\alpha {\rm x}+K}{2K} \tag{57}$$

The last equality is because that $\phi = \frac{\sqrt{(\alpha {\rm x}-K)^2+4\alpha KN}+\alpha {\rm x}-K}{2K}$ is the positive root of ϕ in $\phi+1 = \alpha(\frac{N}{K\phi}+{\rm x})$. ■

Theorem 6: Consider 1-DIA. Suppose $\rho = \frac{1}{2}$ and there is one system having a communication cost per report as 1, while other $(N-1)$ systems having a communication cost per report as $\alpha \ge 1$. There exists no deterministic distributed online algorithm with no intercommunication to achieve a competitive ratio lower than $\alpha(N-1)+\frac{3}{2}$.

Proof: It suffices to consider 1-sDIA. First, we consider a threshold-based algorithm $\mathcal{A}$, such that each system i uses threshold θ_i. It will be generalized to general algorithms.

Let the system having a communication cost per report as 1 be i°. Consider an input σ with two events $\{j_1, j_2\}$ of the following setting:

$$\mathtt{t}^{j_1} = 0, \quad \mathtt{t}^{j_2} = 1, \quad w_i^{j_1} = w_i^{j_2} = \begin{cases} \theta_i+\epsilon, & \text{if } i = i^\circ \\ \alpha\theta_i+\epsilon, & \text{if } i \ne i^\circ \end{cases}$$

for an arbitrarily small ϵ (e.g., $\epsilon = O(\frac{1}{N})$). With no intercommunication, $\mathcal{A}$ needs to produce two report messages from each system, incurring a latency penalty of at least θ_i for the first event for i° and $\alpha\theta_i$ for other $(N-1)$ systems. Note that $\alpha\theta_i \ge \theta_i$. Hence, it follows that

$$\texttt{Cost}(\mathcal{A}[\sigma]) \ge \frac{1}{2}\Big(2\alpha(N-1)+2+\sum_{i\in\mathcal{N}}\theta_i\Big) \tag{58}$$

Recall that $\texttt{Opt}(\sigma)$ is the cost of a (centralized) optimal offline solution, which knows all the inputs in advance without the need for intercommunication. Note that $\texttt{Opt}(\sigma) = \frac{1}{2}\min\{2, 1+\sum_{i\in\mathcal{N}}\theta_i\}$, which either produces two report messages, or just one report message immediately following the second event plus the total latency penalty $\sum_{i\in\mathcal{N}}\theta_i$ for the first event. Hence, the competitive ratio becomes

$$\texttt{CR}(\mathcal{A}) \ge \frac{2\alpha(N-1)+2+\sum_{i\in\mathcal{N}}\theta_i}{\min\{2, 1+\sum_{i\in\mathcal{N}}\theta_i\}} \tag{59}$$

Note that the denominator is $1 \le \min\{2, 1+\sum_{i\in\mathcal{N}}\theta_i\} \le 2$. $\texttt{CR}(\mathcal{A})$ is minimized, when $\sum_{i\in\mathcal{N}}\theta_i = 1$ because

$$x \ge y \Rightarrow \frac{z+x}{1+x} \le \frac{z+y}{1+y}, \text{ for all } x \le 1, z \ge 1 \tag{60}$$

Hence, it follows that $\texttt{CR}(\mathcal{A}) \ge \alpha(N-1)+\frac{3}{2}$.

Next, we consider non-threshold-based algorithm $\mathcal{A}'$. In this case, $\mathcal{A}'$ may have time-varying behavior. But we assume that the behavior of $\mathcal{A}'$ is always identical after resetting $t=0$, because $\mathcal{A}'$ is deterministic. Note that the optimal offline cost $\texttt{Opt}(\sigma)$ does not change. We treat $\mathcal{A}'$ as a blackbox. For each system i, we input a measurement value $w_i = 1$ at time $t = 0$. Suppose that the system waits for τ_i to produce a report message. The competitive ratio of $\mathcal{A}'$ can be lower bounded by the one of a threshold-based algorithm $\mathcal{A}$ with threshold $\theta_i = \tau_i$ using the preceding argument. ■

Theorem 7: For 1-DIA and $\rho = \frac{1}{2}$, there exists no deterministic distributed online algorithm with collision-free intercommunication to achieve a competitive ratio lower than $\frac{\sqrt{N}}{4}$.

Proof: **Step 1**: It suffices to consider 1-sDIA. First, we consider a threshold-based algorithm $\mathcal{A}$, such that each system i uses threshold θ_i. Let $\theta_{\min} = \min_{i\in\mathcal{N}}\{\theta_i\}$. We then generalize the proof to general algorithms. Note that it suffices to consider the case when $\frac{\sqrt{N}}{2}$ is an integer.

First, we construct the two following deterministic inputs:

- $\sigma^{\rm I}$: Every system observes the same measurement values all the time. There are two events $\mathcal{M} = \{j_1, j_2\}$ with the following setting:

$$w_i^{j_1} = w_i^{j_2} = \theta_i+\epsilon, \quad \mathtt{t}^{j_1} = 0, \quad \mathtt{t}^{j_2} = 1 \tag{61}$$

  for an arbitrarily small ϵ.
- $\sigma^{\rm II}$: A particular system is selected by the adversarial to observe non-zero measurement values, while the other systems observe only zero measurement values. There are $\frac{\sqrt{N}}{2}$ events, $\mathcal{M} = \{j_1, j_2, \ldots, j_{\frac{\sqrt{N}}{2}}\}$. Without loss of generality, we assume that system i' has the minimum

threshold such that $\theta_{i'} = \theta_{\min}$, which will observe non-zero measurement values as follows:

$$w_{i'}^{j_h} = \frac{2^h}{\sqrt{N}} + \epsilon, \quad \mathtt{t}^{j_h} = 1 - \frac{1}{2^h} \tag{62}$$

for an arbitrarily small ϵ. And $w_i^{j_h} = 0$ for $i \neq i'$. See Fig. 13 for an illustration.

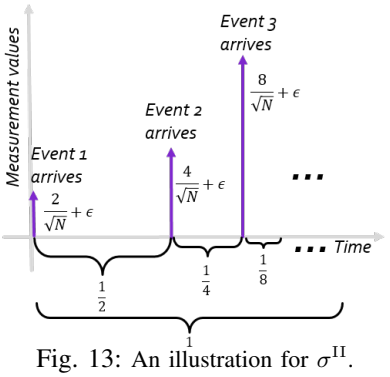

Fig. 13: An illustration for σ^{II}.

If $\theta_{\min} \geq \frac{1}{\sqrt{N}}$, then the adversary chooses input σ^{I}. Otherwise, the adversary chooses input σ^{II}.

For σ^{I}, $\mathtt{Cost}(\mathcal{A}[\sigma^{\mathrm{I}}]) \geq \frac{1}{2}(2 + \frac{N}{\sqrt{N}})$. But the optimal offline cost $\mathtt{Opt}(\sigma^{\mathrm{I}}) \leq \frac{1}{2} \cdot 2$. Therefore, $\mathtt{CR}(\mathcal{A}) \geq \frac{\sqrt{N}}{2} + 1$.

For σ^{II}, only system i' observes non-zero measurement values, which produces a report message for every event (since its threshold $\theta_{i'} < \frac{1}{\sqrt{N}}$). We obtain $\mathtt{Cost}(\mathcal{A}[\sigma^{\mathrm{II}}]) \geq \frac{1}{2}\frac{\sqrt{N}}{2}$. But $\mathtt{Opt}(\sigma^{\mathrm{II}}) \leq \frac{1}{2} \cdot 2$, because one report message is sufficient to be produced at the last event, incurring the total latency penalty $(\sum_{h=1}^{\frac{\sqrt{N}}{2}} \frac{2^h}{\sqrt{N}} \frac{1}{2^{h-1}}) = 1)$. Hence, $\mathtt{CR}(\mathcal{A}) \geq \frac{\sqrt{N}}{4}$.

Step 2: Next, we consider non-threshold-based algorithm $\mathcal{A}'$. Note that the behavior of $\mathcal{A}'$ is always identical after re-setting $t = 0$, because $\mathcal{A}'$ is deterministic. The optimal offline cost $\mathtt{Opt}(\sigma)$ does not change. We treat $\mathcal{A}'$ as a blackbox. First, we input a measurement value $w_1 = \frac{2}{\sqrt{N}}$ at time $t = 0$ to all systems. Suppose that there is intercommunication (possibly, a report message) produced by system i_1 at $t = \tau_1$. Note that any intercommunication among systems is as costly as producing a report message. If $\tau_1 \geq \frac{1}{\sqrt{N}}$, then we apply the argument using σ^{I}, as in threshold-based algorithm $\mathcal{A}$, and obtain the competitive ratio $\mathtt{CR}(\mathcal{A}') \geq \frac{\sqrt{N}}{2} + 1$.

Otherwise, $\tau_1 \leq \frac{1}{\sqrt{N}}$. We input a measurement value $w_1 = \frac{2}{\sqrt{N}}$ at time $t = 0$ to only system i_1. Since $\mathcal{A}'$ is deterministic, there is still intercommunication produced by system i_1 at least at $t = \tau_1$. Then we input a measurement value $w_2 = \frac{4}{\sqrt{N}}$ at time $t = \tau_1 + \epsilon$ to all systems. Suppose that there is intercommunication (possibly, a report message) produced by system i_2 at $t = \tau_1 + \epsilon + \tau_2$. If $\tau_2 \geq \frac{1}{\sqrt{N}}$, then we apply the argument using σ^{I}. Note that the optimal offline cost is less than 2, because one report message is sufficient to be produced at last, incurring the total latency penalty at most 1. We obtain the competitive ratio $\mathtt{CR}(\mathcal{A}') \geq \frac{\sqrt{N}}{2} + 1$.

Iteratively, we apply the same argument up to h times. If $\tau_h \leq \frac{1}{\sqrt{N}}$. We input measurement values $(w_1, ..., w_{h+1})$, where $(w_1, ..., w_h)$ follows from the first h measurement values of σ^{II} at systems $(i_1, ..., i_h)$ accordingly, and we set $w_{h+1} = \frac{2^{h+1}}{\sqrt{N}}$ at time $t = \sum_{h' \leq h} \tau_{h'} + h\epsilon$ to all systems. Suppose that there is intercommunication (possibly, a report message) produced by system i_{h+1} at $t = \sum_{h' \leq h} \tau_{h'} + h\epsilon + \tau_{h+1}$. If $\tau_{h+1} \geq \frac{1}{\sqrt{N}}$, then we apply the argument using σ^{I}.

Finally, if $\tau_h \leq \frac{1}{\sqrt{N}}$ for all $h \leq \frac{\sqrt{N}}{2}$, then we apply the argument using σ^{II}, as in threshold-based algorithm $\mathcal{A}$, and obtain the competitive ratio $\mathtt{CR}(\mathcal{A}') \geq \frac{\sqrt{N}}{4}$. ■

B. Optimal Offline Solution: A Lower Bound

We generalize the offline algorithm in [6] for the 1-report problem to the setting with arbitrary non-decreasing communication cost and latency penalty functions. We denote an offline algorithm by $\mathcal{A}_{\mathrm{ofl}}$ which can produce a lower bound to total cost of the optimal offline solution. The basic idea of $\mathcal{A}_{\mathrm{ofl}}$ is explained as follows. First, a 2D table is created denoted by $\mathtt{Cost}[\cdot, \cdot]$, indexed by the set of events $\mathcal{M} = \{1, ..., m\}$. Each $\mathtt{Cost}[j, \ell]$ denotes the cost when the last report message is produced immediately after event j, and the next report message is produced immediately after event ℓ. $\mathtt{Cost}_{\min}[j]$ is the minimum cost obtained up to event j.

Algorithm 4 $\mathcal{A}_{\mathrm{ofl}}(K, \sigma)$

Input: Number of reports per event K; sequence of measurements $\sigma = (\mathtt{t}^j, (w_i^j)_{i \in \mathcal{N}})_{j \in \mathcal{M}}$

Let $\mathcal{M} = \{1, ..., m\}$
$\mathtt{Cost}[1, 1] \leftarrow \rho \cdot K \cdot \min_{i \in \mathcal{N}}\{\mathcal{C}(\{\sigma_i^1\})\}$
$\mathtt{Cost}_{\min}[1] \leftarrow \mathtt{Cost}[1, 1]$
for $j = 2, ..., m$ **do**
    $\mathtt{Cost}_{\min}[j] \leftarrow \infty$
end for
▷ *Construct a table for dynamic programming*
for $j = 2, ..., m$ **do**
    for $\ell = 1, ..., j$ **do**
        ▷ *Compute latency penalty*
        $\mathtt{lat} \leftarrow \sum_{r: j-\ell < r \leq j} \sum_{i \in \mathcal{N}} \mathcal{L}(\sigma_i^r, \mathtt{t}^r)$
        ▷ *Compute lower bound of communication cost*
        $\mathtt{com} \leftarrow K \cdot \min_{i \in \mathcal{N}}\{\mathcal{C}(\{\sigma_i^r : j - \ell < r \leq j\})\}$

        $\mathtt{Cost}[j, \ell] \leftarrow \rho \cdot \mathtt{com} + (1 - \rho) \cdot \mathtt{lat} + \mathtt{Cost}_{\min}[j - \ell]$

        if $\mathtt{Cost}[j, \ell] < \mathtt{Cost}_{\min}[j]$ **then**
            $\mathtt{Cost}_{\min}[j] \leftarrow \mathtt{Cost}[j, \ell]$
        end if
    end for
end for
return $\mathtt{Cost}_{\min}[m]$

Theorem 8: Algorithm $\mathcal{A}_{\mathrm{ofl}}$ outputs a lower bound to the total cost of the optimal offline solution.

For the 1-report problem, $\mathcal{A}_{\mathrm{ofl}}$ actually outputs the total cost of the optimal offline solution by dynamic programming. For the K-report problem, we obtain the lower bound by substituting the communication cost function by $K \cdot \mathcal{C}(\cdot)$. $\mathcal{A}_{\mathrm{ofl}}$ proceeds to compute the solution as in the 1-report problem. This gives a solution of lower cost than the optimal offline solution, because a report message is replicated K times from the same system has a lower communication cost than the one with K report messages from different systems. The rest of proof is similar to that in [6].